

Creativity and Legitimacy in Artificially Generated Art

A Sociotechnical Research Paper
presented to the faculty of the
School of Engineering and Applied Science
University of Virginia

by

Jeremiah Thomas

May 10, 2021

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

Jeremiah Thomas

Sociotechnical advisor: Peter Norton, Department of Engineering and Society

Creativity and Legitimacy in Artificially Generated Art

In 2018, Christie's art specialist Richard Lloyd organized the \$432,500 auction of a portrait that he called "exactly the kind of artwork we have been selling for 250 years." Lloyd's casual description might strike some traditional painters as strange – the piece was created by a computer and signed with a mathematical equation (Christie's, 2018). Since the early 20th century, advances in technology have challenged traditional notions of art and creativity. During the second half of the century, painters and musicians used computers to blur the lines between artist and tool. Recent breakthroughs in artificial intelligence have led tech enthusiasts to promote the legitimacy of AI-generated art. Increasingly powerful computational tools are gaining acceptance among artists and art critics, but they have mixed feelings. Whether creativity resides in computers or the artists that program them depends on the definitions of art and creativity, and it is an inherently ambiguous question. Some see AI as a more powerful tool and artistic collaborator, while others contend that computers will soon be artists themselves. Previous artistic programs relied on humans, but new artists seek to transcend human intervention.

Review of Research

Harold Cohen's paintings in the 1960s "involved inventing formal elements – things that didn't exist in the real world – and then rendering them as if they did." After finding that inventing new elements was getting harder, he decided to instead "invent a set of rules for making the paintings, so that [he] could then simply follow the rules and not need to invent anything else" (2006). By the 1970s, Cohen had started to turn these rules into AARON, a computer drawing program that functioned as an "extended mind" or "designer environment" that encoded and amplified his knowledge about the rules of color, stroke, and shape without

independent development or creativity (Sundararajan, 2014). Composer David Cope had a similar idea to “code the rules of basic part-writing” for classical music in the 1980s. He called this project Experiments in Musical Intelligence (EMI) and successfully produced “music which basically adhered to these rules.” However, he believed that the “most of the music was uninteresting and unsatisfying” and that the “abstractions of the rule-making and rule-applying processes [had] neutralized any real musical interest.” To remedy this, he “created a new program to derive rules from music coded in a database” (Cope, 2001, p. 93). Cope’s idea of moving rule generation from a human task to a computer task has been explored by researchers working on generative adversarial networks (GANs). In a GAN, the computer learns how to replicate input data by competing with itself. For example, when given a dataset of images, one computer network creates forgeries of the images and another network tries to detect them. As the network evaluating images gets better at detecting forgeries, the network producing images gets better at creating them (Goodfellow et al., 2014). While GANs still form what Cope calls “abstract sets of rules for composition,” they bypass the human “intermediary” that he identified and learn directly from patterns in data (2001, p. 93). Cope relied on a “multifaceted approach to structural analysis” and instructed the program to look for general musical characteristics like “contrast, changes in density, and shifts of composite rhythm” (2001, p. 127). In contrast, GANs receive unstructured data and can produce images without explicit human instructions or input apart from data selection and cleaning.

Advances in technology challenge traditional notions of art and creativity. It was a GAN-generated painting created by French art collective Obvious that Christie’s sold in 2018, and GAN researchers at Rutgers comment that humans were not able to distinguish their AI-produced art from contemporary abstract art (Elgammal et al., 2017). Artists have a variety of

views about machine learning, ranging from treating it as a tool to classifying the trained network itself as a work of art. Mazzone and Elgammal point out that recently, “understanding of art has been expanded to include objects that are not necessarily aesthetic in their purpose (for example, conceptual art), and not created physical objects (performance art)” and suggest that the push for artificial art as legitimate parallels the push for acceptance of photography as art during the 20th century (2019).

Machine Learning as an Artistic Tool

Many tech-savvy artists view AI as an artistic tool. Mario Klingemann, whose GAN-powered art installation, *Memories of Passerby I*, was auctioned by Sotheby’s in 2019, compares his artistic process to “work with brushes or real material” and calls the computer a tool to “help [him] create interesting images” (2018). Similarly, Robbie Barrat, who originally wrote and released the code used by Obvious, views himself as a “curator” and uses GANs as “a tool, strictly” (Bailey & Barrat, 2019). Barrat releasing his code highlights that he views it as a tool for artists to use, and Obvious view it the same way. In response to implications that they would try to patent their algorithm, they clarified that “an algorithm... [is] part of the general knowledge of humankind, and anybody could call it their own and use it” (Bailey, 2018a). Cristóbal Valenzuela, the founder of RunwayML, a machine learning platform that makes programs like GANs more accessible to artists, agrees that machine learning algorithms are “becoming commodities” and that it is the user controlling components of the model that is important. He comments that while some people may think that AI can replace traditional artists, it is actually “just another technology that will help you do things in a better way and express you better” like Photoshop or LightRoom (MIT Docubase, n.d.).

While many artists describe artificial intelligence as a tool, it is clearly a new sort of tool. Unlike Klingemann, Mazzone and Elgammal contend that AI is “more than a tool, like a brush with oil paint on it, which is an inanimate and unchanging object” (2019). Valenzuela believes that artists are moving beyond simply using machine learning technology as a novelty and are “acknowledging the fundamental advantage of using algorithms in their process.” In particular, he says that artists can collaborate with the computer to produce art in less time (MIT Docubase, n.d.). Describing the computer as a collaborator rather than an inanimate object is a common sentiment. Obvious (2020) envision “new generations of creators... where the hand of the artist and the one of the machine are joined,” and Barrat calls his relationship with GANs an “intimate artistic process” (Bailey & Barrat, 2019). Memo Akten, an artist and AI researcher, sees the potential for deep learning to produce visual art in a “realtime, interactive manner, with continuous control – analogous to playing a musical instrument.” In contrast to Valenzuela’s description of RunwayML as a machine learning version of Photoshop, Akten describes his goal as a tool that “becomes an extension of the body, and the act of playing becomes an emotional act in itself” (Bailey, 2018b). The power of machine learning for art and its collaborative nature raises the question of whether creativity in the artistic process resides within the human or the machine.

Toward Creativity

Scientists have long been skeptical that computers could exhibit creativity. In 1949, Geoffrey Jefferson, a British neurologist, commented that “although electronic apparatus can probably parallel some of the simpler activities of nerve and spinal cord,” their “lack of opinions, of creative thinking in verbal concepts” made them more like a crane or an automobile – capable of performing work faster than a human, but not capable of matching the human mind. One of

his criteria for creative thinking was art: “Not until a machine can write a sonnet or compose a concerto, because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain” (Jefferson, 1949). In 1979, shortly before David Cope began to work on EMI, cognitive scientist Douglas Hofstadter echoed Jefferson’s sentiment in *Gödel, Escher, Bach*. While he thought that “shallow imitations of the syntax of earlier music” would be possible in the future, he decided that for a program to produce meaningful music, it “would have to wander around the world on its own... feeling every moment of it. It would have to understand the joy and loneliness of a chilly night wind” (Hofstadter, 1979, pp. 704-705). Pioneering computational artists like Harold Cohen were skeptical, too. While AARON “has a degree of autonomy,” Cohen “never made such a claim” that it was creative in the way that humans are because it lacked “lifelong intellectual development” (2006).

Because computer programs like GANs can now learn from themselves, some claim that they can demonstrate creativity. Jefferson said that computers are limited by the need for “very intelligent staffs to feed them the right problems” and that they will “attempt the insoluble and continue at it” (1949). GANs bypass the first limitation by trying to solve a very general problem (replicate training data) instead of relying on humans to tell them what steps to take to solve a specific problem. However, this does not mean that they are devoid of human input. GANs do teach themselves to replicate data once the program is started, but Akten notes that “[w]hat actually shapes the output and defines what the resulting images look like is the data... and the algorithm,” which are both subject to human discretion (Bailey, 2018b). Barrat agrees “the machine does not really have any agency, it's not doing anything creative” and explains that his “intimate artistic process... with this network” – selecting data, setting training parameters, and curating results – is a necessary part of the method (Bailey & Barrat, 2019). The second

limitation is harder to define and circular. Whether or not creating art is an insoluble problem depends on the definition of art and whether or not one thinks that computers have the ability to create it.

So What Is Art?

The definition of art is ambiguous and fluid, and that is reflected in the views of computational artists. Cope, reflecting on EMI and AARON, believes “without reservation... computer-created art in any form, while it must pass my own personal muster of quality, falls clearly within the same frame of reference as does human-created art.” Outside the frame of human-created art, Cope also thinks that sounds like wind and thunder can count as music (2001, p. 86). Hofstadter on the other hand defines music as something “produced deliberately by a human being in order to express or communicate something to other humans” (Hofstadter qtd. in Cope, 2001, p. 302). Aside from excluding computers from being artists as they are not humans, Hofstadter requires art to have intention and intrinsic meaning. Richard Lloyd of Christie’s takes a pragmatic view, stating that the GAN portrait he sold “is a portrait, after all” (Christie’s, 2018). The range of opinions on computer generated art makes sense given the inherent difficulty in defining art. For some, the difficulty of defining art is the point of making it. Obvious “wish to contribute to the debate regarding the scope and nature of art, and allow once again the definition of art to grow and evolve with the era it evolves in” (2020).

Finding a “true” definition of art to evaluate the creativity of computer-generated art is outside the scope of this paper. However, Pamela McCorduck, who wrote a retrospective of AARON, proposes that because of Cohen’s “intimate and sustained dialogue” with his computer, he has created the first “representation of an artistic *process*, as distinct from an artistic *object*” (1991, p. 195). Whether or not a program is able to possess and express its own creativity,

perhaps it is able to contain creativity from its programmers and reference artists and inspire creativity in observers. Barrat is more explicit in his beliefs and says that “the actual artwork is the trained GAN itself” and calls the individual images it produces “just ‘tokens’ of something larger - they're shadows of the network, the actual artwork” (Bailey, 2019). Mazzone and Elgammal agree that the images are not the full representation of what is going on. They say that AI can be conceptualized as a medium that includes images, tools, critical languages, and “the range of possibilities and limitations inherent to the conditions of creation in that area of art” (2019). So, if there is creativity somewhere – putting aside whether or not the computer is consciously generating meaning – the natural questions are where does the creativity come from, and where does it reside?

Creativity and the Artist’s Role

Many artists and researchers define creativity as turning experience and emotion into something new. For example, Harold Cohen draws a distinction between “the implementation of creativity from its essential core.” Cohen considers himself creative for creating AARON, and concedes that AARON exhibits some components of creativity, but argues “that the essential core of creativity lies, not in its implementation, but in the lifelong intellectual development of the individual and in the highly differentiated world model to which it gives rise” (2006). Hofstadter comments that “there is much more to musical expression than can be captured in syntactical rules” and that “music is a language of emotions, and until programs have emotions as complex as ours, there is no way a program will write anything beautiful” (1979, p.704). Jefferson alludes to this as well, calling a computer no better than “a cleverer parrot” and “typewriting monkeys” unless it can “create concepts and *find for itself* suitable words in which to express additions of knowledge” (1949). For Cohen, Hofstadter, and Jefferson, true creativity

requires continuous learning, emotional capacity, and novelty. The researchers at Rutgers argue that their modified GAN learns from experience and can not only replicate training images, but also “deviate from established styles” provided in the training images (Elgammal et al., 2017). With experience and novelty covered, the missing component is emotion. In his review of EMI, Douglas Hofstadter reflects that he is “not quite sure” if his earlier statements about music and meaning are still true. Even if the program is not capable of feeling emotions, it is still capable of expressing something that is “not emotionally empty” (Hofstadter qtd. in Cope, 2001, pp. 37-38).

Another way to judge creativity is through the interpretation of style. Hofstadter comments that “just a minimal gesture in the direction of a known style can, if well executed, have a stunning effect” (Hofstadter qtd. in Cope, 2001, p. 56). Cope notes that his initial attempt at programming musical rules fell flat because he produced “styleless music that basically adhered to these rules” and that a database would allow a program to assimilate, replicate, and recombine many styles (2001, p. 93). This idea has carried through to AI art. Mazzone and Elgammal claim that the reason their program is creative is because it “simulates how artists digest prior art works until, at some point, they break out of established styles and create new styles” (2019). Other GANs have incorporated style as well, most notably StyleGAN. Using human faces as training data, the authors define style as changes that have “global effects (changing pose, identity, etc.)” as opposed to “inconsequential stochastic variations (differently combed hair, beard, etc.)” (Karras et al., 2019). Faces created by StyleGAN can be seamlessly blended and modified, indicating that the network is not simply copying input images but has acquired a deep understanding of the structure of human faces.

Because GANs still rely on limited human input, some believe that independent creativity will be the result of automation. Commenting on “the debate over whether the computer

simulation of specific musical styles is a machine activity or a human activity,” music researcher Eleanor Selfridge-Field concludes that despite EMI becoming “increasingly automatic,” it still relied on a “great deal of human activity and human judgment.” Because of Cope’s decisions and the inherent structure in composed music, EMI is “generative without being directly creative” and a “reprocessing engine” (Selfridge-Field qtd. in Cope, 2001, pp 215-216). The researchers at Rutgers claim that their approach to automation does not suffer from these problems. While their program cannot spontaneously create art without exposure to existing art, neither can a human artist. Deep learning on an unstructured database of images is “inherently creative” and replicates “the process of how an artist digests art history, with no special selection of genres or styles.” Because of the lack of curation and direction from human artists, they “posit that the person(s) setting up the process designs a conceptual and algorithmic framework, but the algorithm is fully at the creative helm when it comes to the elements and the principles of the art it creates” (Mazzone & Elgammal, 2019). Obvious, while still currently involved in the creative process, agree that “once the whole process will have been automated, we will have created a machine that is capable of being creative, in the same way a human is” (2020).

Conclusion

Contemporary artists were using technology to create art before machine learning, but they viewed “the computer as a rather conventional art-making tool.” Cohen’s creation of AARON, “a program that makes art, not a program with which [he] make[s] art,” signaled a shift from technology as an assistant to technology as a collaborator (2006). In the 20th century, scientists, doctors, and artists were skeptical that computers would ever exhibit creativity. With the rise of deep learning and the increased automation of artistic activities, some say that AI has the potential to be creative, and others say that it already is. Mazzone and Elgammal suggest that

even if artificial paintings are only a reproduction of training paintings, they are still a form of representational art (2019), and programmers like Barrat consider machine learning programs themselves as new form of art. The problem of the possibility of creative AI art is ultimately unanswerable because it depends not only upon the capacities of AI, but also upon divergent philosophical conceptions of what art and creativity are. Who among the computer, the programmer, the computational artist, and the artist in the training data exhibits creativity is still open for debate. As Cohen puts it, “In short, credit assignment is not straightforward” (2009).

References

- Bailey, J. (2018a). The AI art at Christie’s is not what you think. Artnome. <https://www.artnome.com/news/2018/10/13/the-ai-art-at-christies-is-not-what-you-think>
- Bailey, J. (2018b). Machine learning art: An interview with Memo Akten. Artnome. <https://www.artnome.com/news/2018/12/13/machine-learning-art-an-interview-with-memo-akten>
- Bailey, J. (2019). AI artist Robbie Barrat and painter Ronan Barrot collaborate on “Infinite Skulls.” Artnome. <https://www.artnome.com/news/2019/1/22/ai-artist-robbie-barrat-and-painter-ronan-barrot-collaborate-on-infinite-skulls>
- Bailey, J. & Barrat, R. [Christies]. (2019). “Art + Tech Summit: The A.I. Revolution” [Video]. YouTube. <https://www.youtube.com/watch?v=NYaY8rU3h8>
- Christie’s. (2018). Is artificial intelligence set to become art’s next medium? <https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx>
- Cohen, H. (2006). Creativity. <http://www.aaronshome.com/aaron/publications/index.html>
- Cohen, H. (2009). The Art of Self-Assembly: the Self-Assembly of Art. <http://www.aaronshome.com/aaron/publications/index.html>
- Cope, D. (2001). *Virtual Music: Computer Synthesis of Musical Style*. The MIT Press.
- Elgammal et al. (2017). CAN: Creative Adversarial Networks Generating “Art” by Learning About Styles and Deviating from Style Norms. <https://arxiv.org/pdf/1706.07068.pdf>. arXiv.

- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: an eternal golden braid*. New York: Basic Books.
- Karras et al. (2019). A Style-Based Generator Architecture for Generative Adversarial Networks. <https://arxiv.org/pdf/1812.04948.pdf>
- Klingemann, M. [Moving Ideas]. (2018). Moving Ideas.arts - Mario Klingemann - Part 01 of 03 [Video]. YouTube. <https://www.youtube.com/watch?v=FjgZAAQzhGo>
- McCorduck, P. (1991). *Aarons's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen*. W. H. Freeman and Company.
- Jefferson, G. (1949). The Mind of Mechanical Man. *The British Medical Journal*, 1(4616), 1105-1110. JSTOR.
- Mazzone, M. & Elgammal, A. (2019). Art, Creativity, and the Potential of Artificial Intelligence. *Arts*, 8(1), article 26. Web of Science.
- MIT Docubase. Interview with Cristóbal Valenzuela, founder of RunwayML. <https://docubase.mit.edu/lab/interviews/interview-with-cristobal-valenzuela-founder-of-runwayml/>
- Obvious. (2020). Artificial Intelligence for Art Manifesto. <http://obvious-art.com/wp-content/uploads/2020/04/MANIFESTO-V2.pdf>
- Sotheby's. (2019). Mario Klingemann: Memories of Passerby I. <https://www.sothebys.com/en/auctions/ecatalogue/2019/contemporary-art-day-auction-119021/lot.109.html>
- Sundararajan, L. (2014). Mind, Machine, and Creativity: An Artist's Perspective. *Journal of Creative Behavior*, 48(2), 136-151. Web of Science.