What Makes a Masterpiece?

by Audrey Ostrom

This paper covers popular AI models available to the public that are capable of generating images. It then delves into whether or not these outputted images can be considered art by theory - and how these models may be developed and used in the future.

What Makes a Masterpiece?

Audrey Ostrom

***Abstract*—As AI advances, its number of applicable uses expands. AI models have been developed to create high-fidelity images, but there's much debate surrounding them. It's contested whether or not these models are producing art by the literal definition proposed by theorists and psychologists. Regardless, these models are impacting society - though their usage in the coming future remains unknown.**

I. In Recent Media

There's a general fear amongst the populace in this day and age that AI will steal our jobs. This fear applies to the art sphere as well.

Articles published on the subject make commentary on how AI has been stealing awards and monetary rewards from artists. The more sensationalized these articles are, the more waves they tend to make. Even major journals in the U.S. like *The Vulture* and *The New York Times* are covering this topic. Some highlights include "An Artwork Made by Artificial Intelligence Just Sold for $400,000. I Am Shocked, Confused, Appalled."[1] and "An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy"[2]. There's a universal idea amongst these articles, even the ones published by seemingly neutral sources: AI is bad and shouldn't be used for artistic endeavors.

II. But What is AI?

AI, or artificial intelligence, isn't well-defined amongst the general public. A deep understanding of what AI is requires deep technical understanding and a background in machine learning. However, we'll constrain this exploration of AI to how it relates to art.

Over the past half-decade, several deep learning models capable of producing images based on a user-given prompt (such as the two versions DALL-E) have been published. Most deep learning models for this specific task are composed of artificial neural networks, which are made up of neurons, synapses, weights, biases, and functions. In summary, these deep learning models are intended to mimic human cognition - hence the naming of these components and we refer to deep learning models as "AI". It's critical to demystify what's happening behind the scenes with these deep learning models. How are these models processing these user-given prompts to produce an appropriate image? These questions must be answered to properly assess the claims of popular media.

III. Generative Adversarial Networks

The majority of the popular deep learning models available to the public that aren't in beta are considered "generative adversarial networks". This name references how this model is trained to make images. This is a very new type of model in the world of machine learning - it was only proposed back in 2014[3].

Before these models are deployed or accessible to the public market, they are trained on terabytes and terabytes of data and images. Majority of the digital-based training images out there come with accompanying text to them[4]. This caption typically describes the image's context, content, and relationships between significant features, which might not be clearly discernable to the model during feature analysis training.

What makes the training and structure of this model is that it is made up of two different sub-models: a generative network and a discriminative network. Given a fixed-length random vector of input training images, the generative network will try to generate new plausible examples based on it. The discriminator network will then try to differentiate the given sample image from the generated examples by estimating the probability that a sample came from the input image dataset rather than the generated images. The training procedure for this model is to eventually minimize the probability

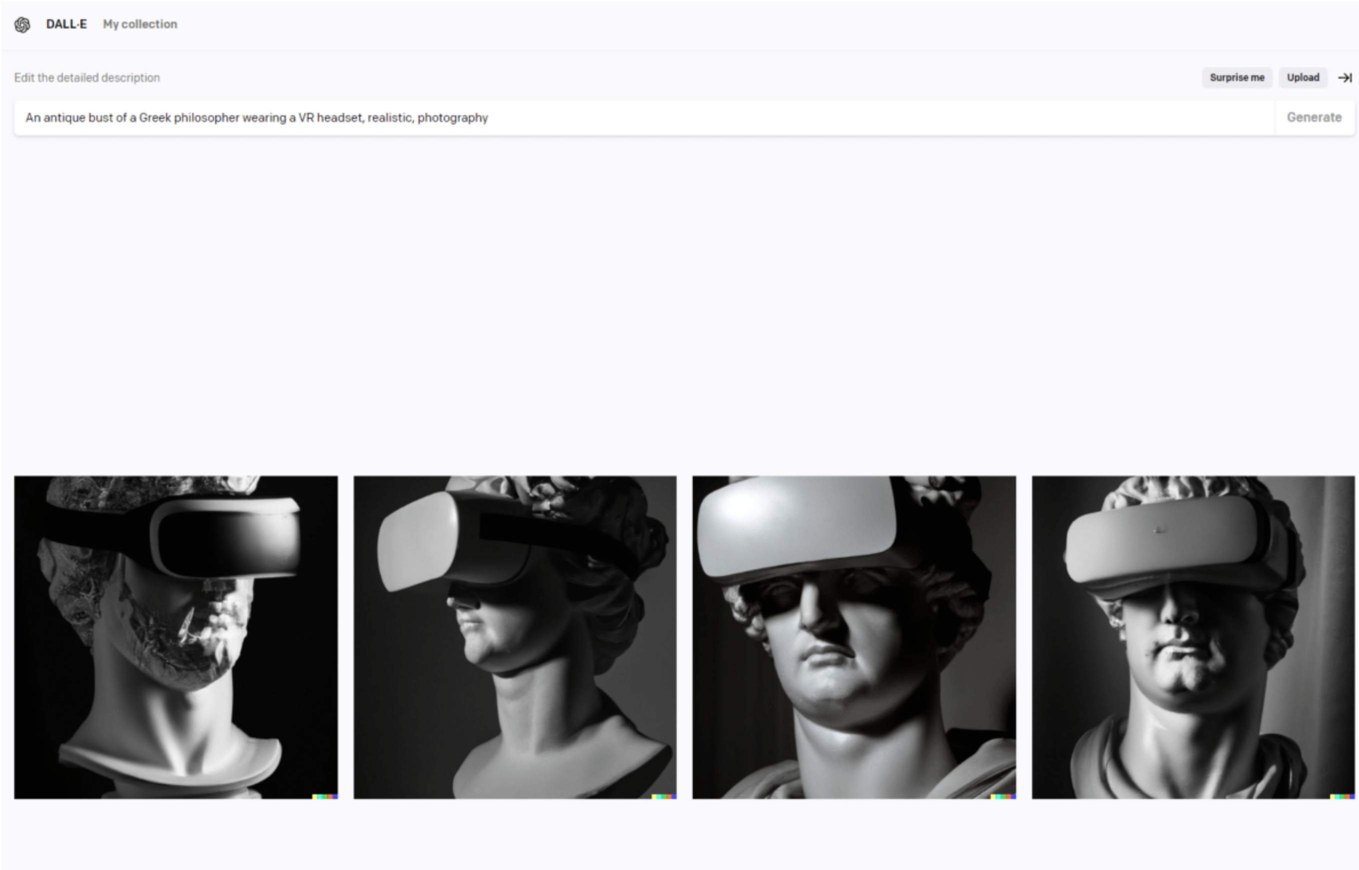

Fig. 1: The interface for DALL-E 2, one of the premier models available

Source: Adapted from [5]

outputted by the discriminator. This training loop will continue until these models have been sufficiently trained for a variety of different possible input vectors.

Most of the models available now have refined this type of model to create images based on user-given prompts. One can use GANs in combination with deep-learning transformer models to create *generative pre-trained transformers* (GPTs). The original version of DALL-E released back in 2021 was built on a generative pre-trained transformer.

IV. From Prompt To Image

After the deployment of the model, the process of generating an image from a prompt can be broken up into stages.

A. Giving a Prompt

For more-developed models, the user will be prompted to give some specification to the image they want to create. Typically, this is done via some text input if the model has a user interface.

B. The Pre-Curation Stage

As mentioned prior, these models are typically exposed for thousands of training epochs on gigabytes of images, so they can fully understand the components in images before deployment. However, when the user gives their prompt to the model after deployment, a fixed-length vector of training images will be selected for image generation later. A query will be sent to find the images whose features and captions best match the prompt in question[**?**]. For example, if the text sub-string “Greek bust” was in the user prompt - it’s highly likely that the majority of the images returned will be photographs of busts from Greece.

C. Evaluation Based on Sample Images

The number of images retrieved in the pre-curation will be of a fixed length, and those will serve as an input to the already-trained general adversarial network. The generative network will generate a series of images, and the discriminative network ensures that the images outputted by the generative network are sufficiently close to the prompt.

D. The Post-Curation Stage

The model will then output a series of images like in Figure 1, and the user will select their favorite amongst them.

V. Agency & Its Importance

Users have the most direct involvement in the pre-curation stages (since their prompt influences image queries) and the post-curation (since they pick the final image to save). However, the model and its database manage everything else - *leaving users with limited agency*.

People place a lot of weight on the artist's agency. Somebody may not fully grasp the meaning or be confused by a cubist painting, but they still value the choices the artist makes - "the intentionality, motivations, and the quality of their work" - even if the work itself may not be distinguishable from a child's painting[6]. It's proposed that's why so many are averse to forgery of art: it takes away from the authenticity and underscores the original creator. Authenticity, then, is the "original conception of the work in the mind of the artist". A forgery is "a refined exercise in paint-by-numbers"[6]. This explains potentially people's aversions to machine-created objects in popular media; the human connection between creator and receiver lends some form of authenticity. The public also deeply appreciates the effort and skill that goes into painting. In some studies, people responded more favorably when shown the same abstract images if told these images were from a gallery rather than computer-generated. There's a sense of community when it comes to art.

On the other hand, some frame art as having no purpose, or "art for art's sake." This sentiment has risen as art becomes more secularized and messages behind art become more opaque. Reproductions are everywhere, widening the disconnect between the original creator and the viewers of the object. With how GANs work, is it apt to describe the outputs of GANs as semi-faithful reproductions of original works? There's potential for AI art to become less stigmatized in the future as this secularization continues.

AI, at the very least, "could be a powerful tool for an artist, perhaps analogous to the way a sophisticated camera is a tool for a fine art photographer"[6]. There's still a human artist dictating the purpose of the art - since their cultural context dictates the prompts they give.

However, AI doesn't learn rules the same way people do, even though the sub-models in the GAN architecture seek to replicate human cognition.

As discussed prior, AI models are fed large sets of data and then trained based on how well they perform a certain function over some sequence of epochs. As such, there are some limits to the originally-proposed architecture of GANs. "They do not possess common sense. They are not adept at analytical reasoning, extracting abstract concepts, understanding metaphors, experiencing emotions, or making inferences"[6]. The model has no agency, which is important for creating art. It is up to the prompt giver to transfer their agency to create true art.

VI. Perceptions of AI Art

Agency is critical to the aesthetics of art - *but does it matter as much if the producer is unknown*? Would we appreciate Picasso's paintings from the blue period if we had no context to the emotions that inspired them? Generative adversarial networks at highly adept at feature replication based on input images, but it's not conclusive whether or not they're as good at replicating the reactions inspired by the original images it's trained on.

In a series of studies performed at the Guangzhou Academy of Fine Arts, they tested willing participants on several art-related topics. In their first study, they asked participants to assess paintings in terms of their "liking rating, purchase intention, and collection intention" without knowing how the painting was made[7]. In their second study, they had art experts assess the same paintings. The authors implicate these results indicate there's "an interaction effect between the author and the art expertise and interaction between the painting style and the art expertise" and that "painting style[s] affects aesthetic evaluation and ultimate reception" - which would impact how AI art may develop in the future.

Recent deep learning networks developed for art generation are trained by learning about styles and how they can be deviated from. Some human subjects are incapable of distinguishing the results of these models from paintings by contemporary artists. However, even with this in mind, the researchers at the academy wanted to explore whether or not AI art captures one's mind. It is one thing to deceive a person into thinking a painting is a masterpiece but it is another to enrapture a viewer entirely. Past researchers on the subject have focused on three main things with AI-generated art: "whether observers could distinguish art generated by AI from those made by humans; whether a bias against AI-created artworks exists; and whether art experience plays a role." In previous studies, they found people can't differentiate between computer-made and human-made art, but there is some bias against AI-generated artwork regardless. People gave these AI pieces a lower aesthetic score generally.

Another important thing to note is that in previous studies: if participants were told that a famous artist made a painting, they would rate it higher than if the same piece was attributed to a lesser-known artist. On the whole, it seems social contexts set the "mental frame that modulates the neurocognitive processing of artworks" - our background dictates our level of engagement with a painting. People have higher neural activation rates when evaluating paintings from the Museum of Modern Art in New York compared to their activation rates when evaluating paintings from an adult education center. We can conclude then that social-identity theory plays a large role in art and its evaluation. There's a sense of ingroup bias in all the evaluations amongst participants in these previous studies. If the participant in question lacked artistic expertise, they would fall back on cultural contexts and identify with what they were familiar with.

It's important to explore people's perceptions of AI-generated art beyond aesthetic evaluations. It's also worthwhile to see if an image being AI-generated affects the purchase and collection intentions of paintings. Working with a team of researchers, they gathered a group of participants with no artistic expertise or background and showed them a series of paintings. Each painting was either human-made or AI-generated and painted in a Western or Chinese style. While it could be easy for even an untrained eye to a painting's style origins, all the information regarding the painting (i.e. who made it, when it was made, etc.) was concealed.

After the first study conducted at the Guangzhou Academy, the researchers found that these non-artists couldn't discern which paintings were AI-generated, and they had no bias against the ones that we indeed AI-generated. However, Chinese participants in this study favored Chinese-style paintings over Western-style ones across all three categories. It's important to remember the studies all took place in China. Continuing from that, most participants preferred AI-generated Chinese-style paintings over AI-generated Western-style paintings.

In their second study, they instead took people with art backgrounds and tested them for the same things. Participants were selected from the Guangzhou Academy of Fine Arts in the design and art education departments, which happens to be the only higher art institution in Southern China approved by the Ministry of Education. Non-art experts also participated in this study. However, their responses exhibited the same general trends as the first study. These art experts evaluated all AI-generated paintings, regardless of their style, lower in all categories (aesthetic rating, purchase intention, and collection intention). However, their ratings across different styles (Western paintings versus Chinese paintings) weren't that different - the results showed no statistically significant differences in their ratings across the three categories.

With these results, the researchers concluded that art expertise and the style of paintings play a role in the assessment of art, be it AI-generated or human-made. One of the greatest influences on the results was ingroup bias, according to them. They also concluded that non-experts favoring Chinese-style paintings and experts having no real preference could be an indicator that "aesthetic judgments" amongst the latter group are "irrespective of [their] cultural background." These results agree with previous studies conducted. However, they did note some other studies found that non-experts rate AI-generated paintings lower, so the conclusions made here may not be universal.

VII. The Theory Behind a Masterpiece

How is it that trained artists are able visually distinct AI paintings from man-made paintings? Is there a secret methodology to classifying what is art?

There's theory to art, just like there is for all crafts. While there are a lot of conflicting ideas as to what constitutes art, there are some universally agreed-upon concepts. One prominent theory is that art "arise[s] from artists' continual necessity to produce novel works in order to counter the effects of boredom or habituation". This need "leads to a monotonic increase across time in the novelty, unpredictability, and complexity of works of art"[8]. In essence, artists will continuously re-invent and redefine stereotypes for art for intellectual engagement. There are several instances of this happening in art history: the development of realism and naturalism in response to the Romantic period of art in the late eighteenth century and the transition from the impressionism era to the post-impressionism in the late nineteenth century are well-documented examples of this.

However, the philosophy behind generative adversarial networks is entirely antithetical to this theory. There's not much room for engineered creativity when the discriminative network is designed to encourage similarity between the generated images and the images from the inputted dataset.

That's not to say there's no purpose in aiming to imitate what already exists. To be a masterpiece, the amount of novelty needs to be minimal to avoid negative reactions from observers. This idea is referred to as the principle of least effort within art circles. New eras of art or style breaks then, under this system of belief, "happen as a way of increasing the arousal potential of art when artists exert other means within the roles of style"[9].

Fig. 2: *Three Studies for a Portrait of Henrietta Moraes* [10]

Fig. 3: Compare the previous figure with these images generated by training a GAN with portraits from the last 500 years of Western art. The distorted faces are the algorithm's attempts to imitate those inputs.

Adapted from [11]

Psychologists focused on the psychology of aesthetics (the study of our interactions with visual arts and culture) have noted the most important properties to the "stimulus relevance to studying aesthetic phenomena" are as follows: "novelty, surprisingness, complexity, ambiguity, and puzzlingness"[12]. But while deformities present in famous paintings like Francis Bacon's *Three Studies for a Portrait of Henrietta Moraes* were intentional to puzzle the viewers, deformities in AI art are decidedly not intentional unless prompted[11]. Machines aren't perfect, so oftentimes we'll see with paintings of human faces that they're jumbled.

GANs simply can't perfectly replicate certain features like a real artist could. However, these deformities still have some aesthetic appeal because of the aforementioned properties of the psychology of aesthetics. Regardless, some artists disregard this machine-created novelty due to a lack of agency and intentionality on behalf of the machine.

Researchers on the subject then propose that's why trained artists can tell an image is AI-generated, when an average person can't. They're trained to have the foreknowledge that no artist in this day and age would try "to emulate the Baroque or Impressionist style, or any traditional style, unless doing so ironically"[9]. If purveying art theory is to be believed, artists would try to increase the arousal potential of their art through novelty, surprisingness, complexity, ambiguity, and puzzlingness.

VIII. How Could We Make These Models Creative?

It's necessary to explore other models in development to answer this question. Several new types of models following the publication of GANs have emerged in very recent years, seeking to refine what originally worked with GANs. The ones that'll be covered here are diffusion models and creative adversarial networks, the former of which is what the second version of DALL-E is based on.

A. Diffusion Models

The major flaw with general adversarial networks for image generation is the lack of originality in outputted images due to the underlying discriminative network. Diffusion models are part of a larger family of "probabilistic generative models" that have overthrown GANs for "dominance...in the challenging task of image synthesis" for that very reason[13]. Instead of relying on a discriminative network to ensure output fidelity, diffusion models will progressively destruct the data or inputted images. They then "learn to reverse this process for sample generation"[13]. Once this process is done for a certain amount of training epochs with a sufficient amount of input images, the model will be able to generate a recreation of the sampled image through the techniques it learned from training on reversing noise.

There are several different techniques for adding and reversing noise in a probabilistic manner. Typically, Markov chains are best for this. A forward chain will be utilized to add noise gradually for image corruption - the goal is typically to transform any data distribution in the sample into a simple prior distribution like the standard Gaussian[13]. A reverse chain will learn to undo the added noise by learning the functions necessary to do so during training. With this process, one can create image that resembles the original sampled image - while not copying it verbatim. This granted flexibility is what makes DALLE-2 so powerful. The image-generating component of DALLE-2 is built on diffusion models developed by OpenAI, one of the biggest pioneers in the field of machine learning at the moment[14]. Several other image generators available to public are also built off of diffusion models: Midjourney, Stable Diffusion, and CLIP-Guided Diffusion are all popular examples.

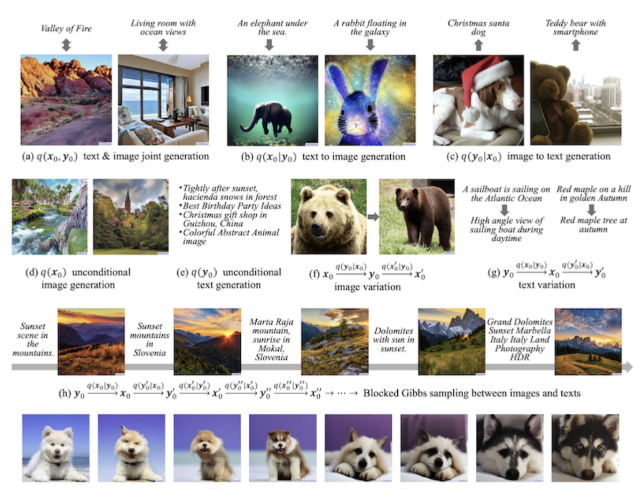

Fig. 4: Diffusion models being used for several tasks, note the variety of outputs.

Source: Adapted from [5]

While user interactivity isn't increased with diffusion models compared to GANs, the fact that diffusion models are capable of producing diverse outputs shows promise for models being used to produce “masterpieces” someday.

B. Creative Adversarial Networks

Currently at the College of Charleston’s Art and Architectural History departments and Rutgers University’s Department of Computer Science, they’re studying how AI processes used for making art can be used to form a partnership between human and machine “creativity” to utilize the strengths of both.

The authors define a term called “algorithmic art” here: which is any art that can’t be created without the use of programming[11]. Art, as a concept, isn’t well-defined. It’s a word used to describe objects that aren’t even intended to be aesthetic (i.e. conceptual art) or physical (i.e. performance art like dancing). However, the meaning piece of “art” is universally dictated by the “determination of the artist’s intention”, its “institutional display”, and “the audience’s acceptance” of the piece[11]. In recent years, the development of GANs (as mentioned prior) has inspired a wave of algorithmic art through the use of AI technologies. The algorithms behind these models are based on learning the aesthetics of art using a generator and discriminator. Creative adversarial networks (CANs) directly build off of the architecture of GANs. However, CANs also use similar techniques as diffusion models. Noise will be continuously added to the inputted images in the same Gaussian fashion, so the generator isn’t as encouraged to “copy”.

Algorithm 1 CAN training algorithm with step size α, using mini-batch SGD for simplicity.

Input: mini-batch images x, matching label $\hat{c}$, number of training batch steps S
for $n = 1$ **to** S **do**
  $z \sim \mathcal{N}(0,1)^Z$ {Draw sample of random noise}
  $\hat{x} \leftarrow G(z)$ {Forward through generator}
  $s_D^r \leftarrow D_r(x)$ {real image, real/fake loss }
  $s_D^c \leftarrow D_c(\hat{c}|x)$ {real image, multi class loss}
  $s_G^f \leftarrow D_r(\hat{x})$ {fake image, real/fake loss}
  $s_G^c \leftarrow \sum_{k=1}^{K} \frac{1}{K} log(p(c_k|\hat{x}) + (1 - \frac{1}{K})(log(p(c_k|\hat{x}))$ {fake image Entropy loss}
  $\mathcal{L}_D \leftarrow \log(s_D^r) + \log(s_D^c) + \log(1 - s_G^f)$
  $D \leftarrow D - \alpha \partial \mathcal{L}_D / \partial D$ {Update discriminator}
  $\mathcal{L}_G \leftarrow \log(s_G^f) - s_G^c$
  $G \leftarrow G - \alpha \partial \mathcal{L}_G / \partial G$ {Update generator}
end for

Fig. 5: The training algorithm of CANs

Source: Adapted from [9]

The user is heavily involved in the pre-curation stage (through the prompts given to the algorithm) and post-curation stage (through their selection of output images).

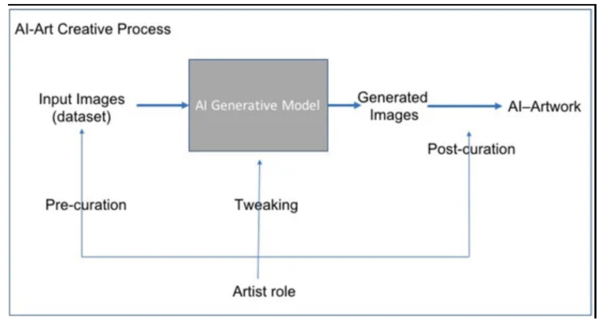

Fig. 6: “A block diagram showing the artist’s role using the AI generative model in making art”

Source: Adapted from [11]

Researchers at Rutgers created a new process using CANs that simulates the way artists digest prior artworks’ style before breaking out and creating their own styles. The architecture of this network is detailed below:

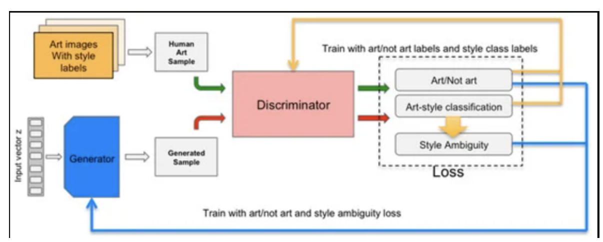

Fig. 7: “A block diagram of a creative adversarial network”

Source: Adapted from [11]

There’s no curation in this dataset - this process is “inherently creative.” This model is learning roughly five centuries of Western art history to better its understanding of art history the way a human artist would. The discriminative network is modified from the GAN structure to discriminate for style rather than a resemblance to a given input image. This will encourage the overall network to output a stylistically creative image.

The rating of this model and its creative functions was based on whether or not people appreciated its outputs as real art. When presenting works created by their creative adversarial network at venues in Los Angeles, New York City, San Francisco, and Miami, the reception of these pieces was overwhelmingly positive amongst people who had no prior knowledge the work was AI-generated. The researchers claim this reception was due to their model's design. It's encouraged not to copy what it's seen, as is typical with GAN-based models, but rather to try out new combinations (due to the creative function of the model).

But are images produced by creative adversarial network models be considered "quality aesthetic experiences?" At the very least, the experiences it produces are more aesthetic than the equivalent experiences produced by GANs using the criteria defined in Section VI through other experiments conducted by the developers to assess paintings on their aesthetic value[9]. However, the proposers of creative adversarial networks and developers of AICAN state their methodology still has its flaws. The outputted images after studying centuries of art styles are too abstract, in their words[9]. They don't look like "traditional art" in terms of standard genres, nor does it seem to portray "recognizable figures"[9]. This is a potential consequence of the encouraged ambiguity.

IX. Advanced, But Not Artistic

There's a general fear amongst the populace against AI-generated art due to the conception that AI will supplant all human artists. The two types of models discussed here (diffusion models and CANs) that build off what worked with GANs, are advanced - but they aren't without flaws. It would be remiss to call these models by themselves artists.

Diffusion models can create high-fidelity images undoubtedly. They can also display some degree of creativity in the sense that the images outputted depart somewhat from the sampled images. However, diffusion models display no personal sense of agency and intention on their own - which somewhat limits their potential unless there's some human intervention.

Creative adversarial networks are indeed computationally creative. They produce stylistically ambiguous images that don't necessarily imitate one specific style, which can be characterized as novel. However, the images outputted tend to be rather featureless. The strengths and weaknesses of diffusion models and CANs contrast each other.

Fig. 8: "Example of images generated by CAN. The generated images vary from simple abstract ones to complex textures and compositions."

Source: Adapted from [9]

At the moment then, AI is too limited currently to be capable of fully creating on its own what has been defined as a "masterpiece": an intentional image that imparts a sense of novelty and imaginativeness that keeps viewers' attention. However, the researcher cites this model can be used as an exploration tool to see what are the "limits of creativity" within the "confines of computation"[9]. Thus, many researchers propose the best course of action is to forge a partnership between man and creative AI systems. AI can be used more as a medium. There was a historied resistance against photographs being considered art when the camera was first released. Now, several art competitions consider photography as its own category: it's art in its own way. Perhaps with more time, AI can be considered the same way a camera or a brush is - a conduit for its user's vision.

X. Looking to the Future

The technical ability of GANs has exponentially increased since their proposal in 2014. We've seen them further developed into even more powerful iterations of the original architecture. AI models for art generation can be used as a powerful tool and the artwork they create can fool most viewers.

However, it hasn't been perfected yet as the models currently don't widely understand all the components to a masterpiece. In the meantime, however, artists can use it to aid their pursuit of art.

It's apparent that the perception of a piece is based on the viewers' biases and culture, but it remains to be seen the ramifications AI will have on this idea due to the recentness of the technology used for art generation. Perhaps models will be promoted to output what would be the most sellable. We are already seeing quite a bit of this in recent media. Photo-realism is in at the moment, and we're seeing models now being specifically engineered to look photo-realistic. When a certain style is favored above all the others, art loses its innovation - and there's no human force with these models to generate novelty unless it's engineered.

Artists such as graphic designers and product illustrators could lose their jobs as AI advances. However, with the nature of humanity, it appears there will be some degree of resistance to this. We as people value agency and intention, so the most ideal scenario for AI in terms of public perception is for it to be a tool. In future models developed, perhaps even more emphasis will be placed on the pre-curation and post-curation stages. Users could have direct control over the databases used for training and could be able to reward certain outputs of the model. As it stands, the future isn't a clear picture as machine learning standards continuously reinvent themselves year by year.

References

[1] J. Saltz, "An Artwork Made by Artificial Intelligence Just Sold for $400,000," *Vulture*, Oct. 25, 2018. https://www.vulture.com/2018/10/an-artificial-intelligence-artwork-just-sold-for-usd400-000.html (accessed Jun. 06, 2023).

[2] K. Roose, "An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy.," *The New York Times*, Sep. 02, 2022. Accessed: Jun. 06, 2023. [Online]. Available: https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html

[3] I. J. Goodfellow *et al.*, "Generative Adversarial Networks." arXiv, Jun. 10, 2014. doi: 10.48550/arXiv.1406.2661.

[4] E. Cetinic and J. She, "Understanding and Creating Art with AI: Review and Outlook," ACM Trans. *Multimedia Comput. Commun. Appl.*, vol. 18, no. 2, p. 66:1-66:22, Feb. 2022, doi: 10.1145/3475799.

[5] "OpenAI DALL-E 2 Prompt Guide: How to control image generation." https://the-decoder.com/openai-dall-e-2-prompt-guide-how-to-control-image-generation/ (accessed Jun. 10, 2023).

[6] A. Chatterjee, "Art in an age of artificial intelligence," *Frontiers in Psychology*, vol. 13, Nov. 2022, doi: 10.3389/fpsyg.2022.1024449.

[7] L. Gu and Y. Li, "Who made the paintings: Artists or artificial intelligence? The effects of identity on liking and purchase intention," *Frontiers in Psychology*, vol. 13, Aug. 2022, doi: 10.3389/fpsyg.2022.941163.

[8] C. Martindale, "Aesthetic evolution," *Poetics*, vol. 15, no. 4, pp. 439–473, Dec. 1986, doi: 10.1016/0304-422X(86)90006-9.

[9] A. Elgammal, B. Liu, M. Elhoseiny, and M. Mazzone, "CAN: Creative Adversarial Networks, Generating 'Art' by Learning About Styles and Deviating from Style Norms." *arXiv*, Jun. 21, 2017. Accessed: Jun. 13, 2023. [Online]. Available: http://arxiv.org/abs/1706.07068

[10] F. Bacon, *Three Studies for a Portrait of Henrietta Moraes*. Private gallery, 1963.

[11] M. Mazzone and A. Elgammal, "Art, Creativity, and the Potential of Artificial Intelligence," *Arts*, vol. 8, no. 1, p. 26, Feb. 2019, doi: 10.3390/arts8010026.

[12] D. E. Berlyne, "Aesthetics and Psychobiology," *Journal of Aesthetics and Art Criticism*, vol. 31, no. 4, pp. 553–553, 1973, doi: 10.2307/429334.

[13] L. Yang *et al.*, "Diffusion Models: A Comprehensive Survey of Methods and Applications." *arXiv*, Mar. 23, 2023. Accessed: Jun. 13, 2023. [Online]. Available: http://arxiv.org/abs/2209.00796

[14] "DALL·E 2." https://openai.com/dall-e-2 (accessed Jun. 13, 2023).