A.I Generated Art Examined in Different Scenarios in Respect to Authorship

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Charles Monsell

Fall 2023

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

Advisor

Kathryn A. Neeley, Associate Professor of STS, Department of Engineering and Society

Introduction – The Growth of A.I Research and Controversies Surrounding A.I Generators

Recent advancements in computing power have led to the meteoric rise of artificial intelligence research and development. Kobe Millet, a researcher from the University of Amsterdam writes "Artificial intelligence (AI) has grown exponentially in the first decades of the 21st century to the extent that our time has been called the age of AI or the fourth industrial revolution" (Millet, pg. 1). Companies such as OpenAI have made waves on the market with their generative AI tools while making over $100 million a month from its subscription-based model (McKay, 2023). These tools enable users to produce creative works with only a few sentences of input that is on par with the works of professionals in some cases.

While A.I is an extremely broad field, one of the most sensitive topics surrounding it is the outcomes of A.I generated art. In this case, Jessica Fjeld, a Harvard researcher defines A.I art in this context as "Advanced algorithms that display cognition-like processes and are capable of generating sophisticated and provocative works of art" (Fjeld, 2017). They are fundamentally different from traditional art tools in the sense that they can make complex decisions which lead to unexpected outcomes (Fjeld, 2017). Beyond definitions, a basic understanding of how these tools generate outputs is required when discussing many of the controversies surrounding it. In 2022, Cetinic and She, engineers from Hong Kong University, describe one of the most widely used methods for generative A.I technology, generative adversarial networks (GANs): "The key mechanism of a GAN is to train two "competing" models that are usually implemented as neural networks: a generator and a discriminator" (Cetinic & She, 2022). The goal of the generator is to capture the distribution of true examples of the input sample and generate realistic images, while the discriminator is trained to classify generated images as fake and the real images from the original sample as real. Designed as a minimax optimization problem, the optimization process

ends at a saddle point that is considered a minimum in relation to the generator and a maximum to the discriminator (Cetinic et all, 2021). Further inspection of the inner workings of this technology reveals that these tools lack creative capacity: “GANs have a very limited capacity for autonomy because they synthesize images that mimic the latent space of the training data, but do not have any role in choosing the input datasets nor the statistical model that represents the latent space. In this sense, GANism is more a process of mimicry than intelligence” (McCormack, pg. 5). Essentially, the key to creating these tools is using massive amounts of data, likely pulled from the work of actual artists, and creating an algorithm to make a prediction on what the output should be based on the minimal input of the user.

Lastly, the notion of authorship in this case drives much of the debate surrounding the legal landscape. Christopher Zirpoli, a Legislative Attorney mentions “Given the lack of judicial or Copyright Office decisions recognizing copyright in AI-created works to date, however, no clear rule has emerged identifying who the “author or authors” of these works could be” (Zirpoli, 2023). From this, I argue that the outputs of these tools are not equivalent to human created works in respect to authorship. To show this, A.I art generators are examined in different scenarios and through different perspectives to show how the language used to describe human art is vastly different from the language reserved for A.I art.

Section 1 – A.I Generated Art vs. Human Created Art Depends on the Definition of Authorship

A.I generated art as it stands currently represents an area of legal greyscale in respect to copyright, data privacy, and other intellectual property laws. It is currently unclear where A.I generated art fits in the legal landscape. This paper examines the language surrounding the current opinions surrounding A.I art generators to determine where they fit in in respect to authorship and the scenarios in which authorship is ill-defined.

Currently, one of the largest controversies surrounding A.I art generators is that training models require extensive amounts of data. Rachel Gordon, a researcher from MIT succinctly illustrates the problems raised by A.I generators in respect to intellectual property: “Since these models are trained on vast swaths of images from the internet, a lot of these images are likely copyrighted. You don't exactly know what the model is retrieving when it's generating new images, so there's a big question of how you can even determine if the model is using copyrighted images. If the model depends, in some sense, on some copyrighted images, are then those new images copyrighted?” (Gordon, 2022). The act of training models directly leads to the problem of defining authorship for A.I generators because the training data largely involves copyrighted work of human artists. This quote also suggests that the outputs of these tools are merely derivatives of the data used to train the models. This raises issues if the data used for training contains copyrighted work of real artists.

From the perspective of artists, Opinions on this technology are extremely divisive with many artists discrediting the effort it takes to develop the prompts to produces the outputs as seen in the article by Kevin Roose who quotes an artist saying “This is so gross, I can see how A.I. art can be beneficial, but claiming you’re an artist by generating one? Absolutely not” (Roose, 2022). Artists generally discredit the lack of effort required to create A.I art pieces citing the lack of creative involvement. Likewise, artists are also upset with the use of their works in training data as seen in an article from Melissa Heikkila, “surprised by the apparent popularity of their work in text-to-image generators—and some are now fighting back. Karla Ortiz, an illustrator based in San Francisco who found her work in Stable Diffusion’s data set, has been raising awareness about the issues around AI art and copyright” (Heikkila, 2022). As it stands now, the

AI's user, the AI's programmer, and the AI program itself all play a role in the creation of these works which makes it hard to determine who gets to take the crown of authorship.

While the ill-defined definition of authorship with respect to this technology muddies the waters when assigning copyright protections, even more complexity is added when these tools are used to mimic the work of other artists. Elizabeth Penava collected accounts from artists whose works, styles, and idiosyncrasies have been mimicked by A.I generators such as one account from classical painter Greg Rutkowski who "has reportedly complained that AI-generated images mimicking his art are drowning out his own work. Users had apparently prompted Stable Diffusion with text including Rutkowski's name nearly a hundred thousand times as of September 2022" (Penava, 2023). Going back to the notion of training and how artists' works are widely used in the training data, they have every right to be upset especially when these tools are used to create art in their styles. To show how easy this is Figure 1 shows an example of how these tools can easily be used to create art in likeness to other artists. The prompt for the tool was "Create a piece similar to Ballon Venus by Jeff Koons".

Figure 1: *Balloon Venus* by Jeff Koons on the Left (Koons, 2017), Piece by Dalle2 with a Prompt from the Author (Created by Dalle2)

From a legal perspective, copyright laws place a lot of emphasis on the notion of authorship as this determines the protections that a work may enjoy (Zirpoli, 2023). A paper about a survey surrounding the legal landscape around this technology writes “Artificial intelligence (AI) and copyright law intersect when copyrighted data are used to train machines to learn, reason, and act as humans do” (Quang, pg. 4). As the development of these models require extensive amounts of data to “train” from, it makes it difficult to assign copyright protections to the outputs of these tools. While it is not well defined what an author is in a legal sense, the U.S. Copyright Office recognizes copyright only in works “created by a human being”. This raises the question if the act of writing the prompt is sufficient human input to warrant copyright protection for the person that wrote the prompt: “Assuming that a copyrightable work requires a human author, works created by humans using generative AI could still be entitled to copyright protection, depending on the nature of human involvement in the creative process” (Zirpoli, 2023). Many of the questions surrounding this controversy hinge on whether the act of writing a prompt for one of these tools is sufficient human involvement. Most people, and courts, ration that these tools bypass the creative process outright (Zirpoli, 2023). This sentiment is also reflected by artists as seen above. Lastly, an article from a law firm examines the use of A.I generators in this fashion where they mention the “personal touch test” which classifies human involvement: “In this scenario, the human authors have a passive role, but they can make creative choices in the creation process. Here, AI generates the work independently without human involvement. It most likely would not meet the personal touch test” (Solyom, pg.1). Clearly these tools facilitate the unfair use of artists’ works.

Lastly, an anthropomorphic perspective is useful in showing that the output of these tools does not share the same creative weight as actual art. The main appeal of art is its ability to serve

as humanistic expression. Using predictive or generative tools largely removes that. In respect to art and anthropomorphism, Kobe Millet et al wrote that "recent advances of artificial intelligence (AI) in the domain of art (e.g., music, painting) pose a profound ontological threat to anthropocentric worldviews because they challenge one of the last frontiers of the human uniqueness narrative: artistic creativity" (Millet et al, 2023). From this, it is clear that people are sensitive to the automation of human expression. Beyond lacking creative expression, Kieran Browne in his 2020 article wrote how an A.I was able to generate a modern contemporary art piece that was then sold for $432,000 (Browne, 2023). This article suggests that the price tag of this art piece solely results from novelty. A.I art is a popular buzzword that will likely decay in value and people will see it for what it is: a generative tool that makes works that copy existing styles. This is further seen in a study produced by Joo-Wha Hong and Nathaniel Curran in which subjects evaluated human and A.I art side by side, without knowing who or what created the works, concluded that "this survey experiment indicate clear differences in evaluation between human created artworks and AI-created artworks, and it is possible to assume that such difference is due to human-created artworks having a significantly higher rating in "composition," "degree of expression," and "aesthetic value"" (Hong and Curran, 2023). Various other studies have concluded similar results such as one seen in "Does Human-AI Collaboration Lead to More Creative Art?" from Jimpei Hitsuwari in that while these tools can produce shockingly well-crafted results, people generally do not view them on equal footing with human work in a creative sense. With this, A.I generators clearly cannot be authors in the same way that humans are. They are unable to generate new ideas and merely mimic existing work which is reflected by its lack of expression.

Section 2 – Utilizing Discourse Analysis to Construct a Narrative on A.I as an Author

As a general overview, discourse analysis provides a methodological framework for examining the role of language in shaping our perceptions of the world and the ways in which it reflects and influences social structures and power dynamics. It's a valuable tool for researchers interested in understanding how communication shapes society and culture (Luo, 2019). Performing a discourse analysis is typically done via a literature review. This entails compiling a basis of textual evidence of the chosen topic. This paper analyzes the language surrounding A.I generators by using many direct quotes from various sources to show why the use of A.I generators to create art is not the same as creating art by hand.

Given that one of the most important questions with respect to this technology is nailing down the definition of authorship, a lot of this paper will focus on looking at how the language surrounding authorship in respect to these tools in various scenarios and perspectives. The goal is to help define the line as to when using tools bypasses the creative process all together. This is important as The United States Copyright Act embraces "original works of authorship fixed in any tangible medium of expression, now known or later developed, from which they can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device" (*17 USC 102: Subject Matter of Copyright: In General*).

A lot of emphasis was placed on the perspective of artists as they are most impacted. For example, in Kevin Roose's article he states accounts from some artists on social media in respect to an A.I generated work that won a competition in which one artist states that "We're watching the death of artistry unfold right before our eyes," (Roose, 2022). Generally, most artists are against this technology as they claim that it bypasses the creative process, and that the training data consists of their work. Specifically, they argue that this tool differs from tools created in the

past because of the act of training which involves compiling mass amounts of data which includes works of the artists who are affected by this.

Lastly, the changes and shifts in power dynamics that these tools may raise were examined by considering the legal implications in respect to copyright and intellectual property. Copyright laws were briefly examined to highlight how current legal opinions are largely against granting the users of these tools copyright protections because of the lack of human involvement.

Section 3 – Perceptions of A.I Authorship are Uniform in Various Scenarios

Something that became apparent from reviewing the variety of sources compiled was that the perceptions of A.I art vary depending on the context. With the information gathered from the literature review, a narrative is constructed on the concept of authorship in respect to A.I generated art. This is used to help show how these tools bypass the creative process. In all the scenarios and perspectives analyzed it was clear that using these tools to "create" art is not equivalent to creating art by hand. This is true for the A.I and the user that prompts the A.I to create the output.

Opinions about A.I generators fell into three categories, legal, anthropomorphic, and technological. The legal opinion corresponds with how this technology falls into the legal landscape, anthropomorphic corresponds to objective comparisons of A.I and human created arts humanistic aspects, and technological corresponds to evaluations of authorship based on how outputs are generated. Looking at these categories, technological and legal are the most important when considering changes to power dynamics as the creation of policy has the most direct impact on how these tools interact with society. However, the humanistic or anthropomorphic side reveals more about human involvement in the creation of art which is the backbone of the legal landscape.

Looking at the legal perspective, there is clear emphasis on humanistic expression as seen in copyright laws and in Zirpoli. In this case, the lack of human involvement when using these tools is the driving force for denying copyright protections to the users who prompt the A.I to generate something. In respect to intellectual property, individual cases in which people use these tools to create art in which they sell are important in showing the harm these tools can do. Browne makes this clear when discussing an A.I art piece that was sold for $432,00. Here, the "artist" of the piece is rather humble about taking credit for the work and can acknowledge that the act of using a generative tool is not the same as creating it themselves. This opinion is then contrasted by the fact that the person who bought this piece was willing to pay an exorbitant amount of money for it. Looking at the art piece, as seen in Figure 2, it does not appear to be objectively better than human made contemporary art which raises the question as to why someone would be willing to spend that much money on it.

Figure 2: A.I Generated Painting, Sells for $400k (Browne, 2020)

Furthermore, the painting does not showcase any distinguishing or innovative techniques or designs. Instead, it mocks the style of already existing art. This is reflected by Penava who shows how people are using these tools to create works that mimic the work of real artists. Given that the training data consists of the work that the generator is mimicking, artists can potentially make the case that their work is being used unfairly. This sentiment is also reinforced based on how these generators work.

Next, the anthropomorphic perspective is a particularly sensitive one. This is because for most of human history, the biggest line that divided man from technology was mans ability to go though the creative process. People have a hard time acknowledging a machine's ability to "create". However, this goes beyond just a simple targeted or pervasive bias against A.I art from people. Hong and Curran showed with their study that people generally do not perceive the expressive properties of A.I art in the same way they do with human created art. In this case, they described the A.I art pieces as "lacking expression" which is very telling considering the subjects in the study had no prior knowledge of who or what created the pieces that they evaluated (Hong and Curran, 2023). The notions of creativity and expression are what drives the debate on determining whether this technology counts as a tool. People, especially artists, are upset that users can essentially create relatively high-quality art instantaneously. They essentially claim that the use of these tools bypasses the creative process all together. On the contrary, users of these tools cite the creative merit required to come up with the prompts to pass in. They argue that this notion of prompt engineering requires at least some degree of savviness and creativity. Likewise, another argument is that the goal post for what is and is not too powerful constantly changes as technology develops. Similar feelings of dismissiveness can be seen with the development of past art tools which decayed over time as the tool became more and more utilized. However, the

new technology argument falls short as the user was still responsible for creating their idea themselves. In this case, the idea is materialized for them. The intricate details, arguably the most important part of a creative work, are filled in by the generator which most definitely cheapens the creative merit in respect to a piece created from scratch. The tools predict what should come next based on the user's prompt. Once again, understanding the mechanisms that drive these generators reinforces the sentiments people express about their outputs.

Lastly, the technological perspective reinforces the sentiments and dismissive opinions people have about its use for artistic expression. Specifically, how the act of training models on copyrighted works blurs the line as their generated outputs fit in in the legal landscape. At a glance, the outputs of these tools seem impressive in terms of creative ability. However, this is simply because the data sets used to train commercial models are so large that they can feign actual creativity. Fundamentally, the outputs are products of the data the model is trained on which is supported when looking at how GANs generate outputs as seen in McCormick. This highlights two things. First, the A.I tools are not capable of being authors in a humanistic sense. Second, a user prompting the tools to do something is not an author either as the tool just generates art with art in its database that most closely resembles the prompt in the first place. From the categories mentioned above, it is clear that A.I art and human art do not share the same perceptions. Even in different circumstances, A.I art and human art are still viewed drastically different. In respect to authorship, this means that none of the stakeholders have the right to claim authorship of the tool's outputs.

Conclusion – A.I Art and Creativity

The usage of A.I tools to create art has proven to be extremely controversial specifically because of the ease of use and the quality of its outputs. The goal of this project was to highlight

how notions of authorship get distorted when using tools that fill in the details when attempting to manifest an idea into an art piece. This directly leads to conflicting arguments when discussing copyright and intellectual property law as the protections granted by these laws hinge on whether the person demanding the rights can be considered an author. Examining authorship in this context led to the discovery that perceptions of this technology are uniform in many scenarios. Generally, the use of A.I generators to create art is not creating art at all. In this case, three distinct perceptions led to this conclusion, legal perceptions, anthropomorphic perceptions, and technological perceptions. Each case supported the fact that these tools do not constitute artistic expression largely because of the lack of human involvement. Furthermore, the predictive nature that the technology is built from means that the work created by this tool will always be a derivative of their training data and that no new styles or ideas can be expressed. The results of this paper highlight how the usage of these tools can potentially stagnate the expression of new ideas. They cannot be used to create new styles or make any innovative contributions; They can only mimic existing ones.

References:

By, W., & Gordon, R. (n.d.). *3 questions: How ai image generators work*. 3 Questions: How AI image generators work | MIT CSAIL. https://www.csail.mit.edu/news/3-questions-how-ai-image-generators-work

E. Cetinic and J. She, "Understanding and Creating Art with AI: Review and Outlook," ACM Trans. Multimed. Comput. Commun. Appl., vol. 18, no. 2, pp. 1–22, May 2022, doi: 10.1145/3475799.

E. Penava, "AI Art Is in Legal Greyscale," Jan. 24, 2023. https://www.theregreview.org/2023/01/24/penava-ai-art-is-in-legal-greyscale/

Heikkil, M. (2022, November 23). This artist is dominating AI-generated art. and he's not happy about it. MIT Technology Review. https://www.technologyreview.com/2022/09/16/1059598/this-artist-is-dominating-ai-generated-art-and-hes-not-happy-about-it/

J. Fjeld, "A Legal Anatomy of AI-generated Art: Part I", [Online]. Available: http://jolt.law.harvard.edu/digest/a-legal-anatomy-of-ai-generated-art-part-i

J.-W. Hong and N. M. Curran, "Artificial Intelligence, Artists, and Art: Attitudes Toward Artwork Produced by Humans vs. Artificial Intelligence," ACM Trans. Multimed. Comput. Commun. Appl., vol. 15, no. 2s, pp. 1–16, Apr. 2019, doi: 10.1145/3326337.

K. Browne, "Who (or What) Is an AI Artist?," Leonardo, vol. 55, no. 2, pp. 130–134, Apr. 2022, doi: 10.1162/leon_a_02092.

K. Millet, F. Buehler, G. Du, and M. D. Kokkoris, "Defending humankind: Anthropocentric bias in the appreciation of AI art," Comput. Hum. Behav., vol. 143, p. 107707, Jun. 2023, doi: 10.1016/j.chb.2023.107707.

K. Roose, "An A.I.-Generated Picture Won an Art Prize. Artists Aren't Happy.," Sep. 2022, [Online]. Available: https://www.nytimes.com/2022/09/02/technology/ai-artificialintelligence-artists.html

Luo, A. (2019, August 23). Discourse Analysis | A Step-by-Step Guide with Examples. Scribbr. https://www.scribbr.com/methodology/discourse-analysis/

McCormack, J., Cruz Gambardella, C., Rajcic, N., Krol, S. J., Llano, M. T., & Yang, M. (2023, April). Is Writing Prompts Really Making Art?. In International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar) (pp. 196-211). Cham: Springer Nature Switzerland.

McKay, C. (2023, October 12). *OpenAI's revenue skyrockets to $1.3 billion annualized rate*. Maginative. https://www.maginative.com/article/openais-revenue-skyrockets-to-1-3-billion-annual-rate/#:~:text=This%20means%20the%20company%20is,from%20just%20this%20past%20summer

Quang, Jenny;, "Does Training AI Violate Copyright Law?," 2021, doi: 10.15779/Z38XW47X3K.

Solyom, A. (2019, July 10). *Ai created works – creative, original or just sweat of the brow?*. CMS Law. https://cms.law/en/media/local/cms-cmno/files/publications/publications/ai-created-works-creative-original-or-just-sweat-off-the-brow

Z, Christopher. U.S. Congressional Research Service. Generative Artificial Intelligence and Copyright Law (LSB10922; May 11, 2023). Text in: ProQuest® Congressional Research Digital Collection.

17 USC 102: Subject matter of copyright: In general. (n.d.). Uscode.house.gov. https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title17-section102&num=0&edition=prelim