

# **Attitudes Towards AI Generated Art**

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

**Andrea Jerausek**

Spring 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

Rider W. Foley, Department of Engineering and Society

## **Motivation Behind Studying AI Art Generation and Opinions of the Public**

When a new AI model is created, controversy is sparked and concerns about the over reliance of technology arise. One source of controversy emerges due to people viewing AI as a “threat or rival” to human labor (Mazzone & Elgammal, 2019). It is feared that since robots have replaced humans in completing menial tasks in factories, they will eventually replace humans for creating art. Artists’ fear of replacement by AI is derived from the high quality of recent AI produced art (Roose, 2022). There is also a concern about whether AI generated art is stealing work from human artists as AI often uses previous artwork for training or stylizing purposes (Salkowitz, 2022). Negative bias toward AI generated art is also derived from the belief that AI is unable to be creative because of its inability to communicate ideas from art to the audience (Yamshchikov & Tikhonov, 2018).

Not only that but according to one study, researchers found that opinions toward AI generated artwork were negatively affected when participants were reminded that the AI does not have a sense of self (Lima et al., 2021). This indicates that people’s ontological analysis of machines may prevent them from even considering art created by AI as true art (Lima et al., 2021). Although there are a lot of negative opinions directed towards AI generated artwork, one study found that people tend to unknowingly rate AI generated artwork about the same as human generated art when considering aspects such as originality, degree of improvement or growth, development of personal style, experimentation or risk taking, and communication of ideas (Hong & Curran, 2019). Since a disparity between opinions towards AI generated art is created when the artist’s identity is transparent, it makes one question whether these negative biases are rational. Because negative perceptions of AI generated artwork might be irrational, it is important to verify whether people’s understanding of AI models and model training will

influence their perception of the technology. Overall, this research studies if people's attitudes towards AI generated art change if they learn about the DALL·E models mechanisms.

## **Basic AI Concepts**

Before defining how the DALL·E model operates, we should establish a basic understanding for the terms AI, models, and training. First, what differentiates general programs from artificial intelligence is that when general programs execute, they either take in an input or execute a unit of code when asked and produce an output that is deterministic. However, with AI, the decision making of this software should mimic that of a human. Therefore, while one can have a set of expectations for the output an AI will produce, it may do something random. This aspect of randomness is similar to the decision-making process of a human. For instance, if Bill goes to get coffee from Starbucks every day before work, we can hypothesize that he will go get coffee from Starbucks tomorrow. However, Bill may choose to be adventurous and go to Dunkin Donuts instead for no clear reason. AI may make similarly random decisions when run.

AI models are concerned with *how* a particular AI will make decisions. On a lower level, the *how*, relates to specific techniques employed that will influence the output the AI will produce. The techniques are heavily dependent on various mathematical concepts. Asking various AI models to complete a specific task is comparable to asking two different people to drive to Starbucks. They'll both probably get in a car and attempt to reach Starbucks (or at least a coffee shop) but the exact route they take to reach their destination may be different.

*Training* AI models directly relates to the input you feed the algorithm. In order to produce answers similar to humans, the algorithm needs to experience information the same way humans do. The training will influence the AI models' outputs greatly. This training is related to how humans are influenced by the environments they grow up in. Someone growing up in rural

Texas versus New York City are going to have different values, personalities, and biases. The initial environment they grow up in will directly influence how each person behaves later on in life. The same is true for AI models. If you train the model on one specific set of inputs, it will later produce outputs that are directly biased towards what they were initially exposed to.

Here, we are particularly interested in understanding how people's perception of art from the DALL·E model is influenced by their understanding of the model's mechanism. The DALL·E model does text to image generation, and it generates art with an autoregressive approach (predicting future outputs based on previous inputs) (Ramesh et al., 2021). DALL·E was trained with 12 billion input parameters and "250 million image-text pairs," where text and images were fed into the model as one stream of data, that were unlabeled images obtained from the internet and no human verified image description (Ramesh et al., 2021).

### **Defining DALL·E using Latour**

Since the analysis of a particular AI model is important to this research, defining the model in a particular way is important with this assessment. In order to analyze issues people might have with the DALL·E model generating art, Latour's definitions of prescription, circumscription, and description will be utilized. Prescription is defined as the "moral and ethical dimension of a mechanism" (Latour, 1992, pp. 157). The paper will use prescriptions to define *what* values AI generated art has and who the technology discriminates against. Circumscription explains the limits the technology has due to external factors (Latour, 1992, pp. 162-163). Here, it will be used to identify *when* art generated by AI affects users. Description of AI generated art defines the exact mechanism that the technology is built on and it will be used to define *how* AI generates art (Latour, 1992).

To start, let us understand what values this technology prescribes. The DALL·E model relies on a massive amount of data that can be accessed for free on the internet. The primary reason this technology exists is because it was able to train on a stream of data that consisted of “250 million image-text pairs” (Ramesh et al., 2021). Since this technology values processing massive amounts of data that is neatly organized on the internet, it also consumes massive amounts of energy. The exact amount of energy that the DALL·E model consumes is not public knowledge; however, it is generally understood that AI creates environmental problems due to its large carbon footprint from both data center energy consumption and the building of super computers needed to run complex AI algorithmic models (Lai et. al., 2022). However, OpenAI, the creator of DALL·E, has been working with Microsoft’s cloud platform Azure (Brockman et. al., 2016). This piece of data is important to further understand the values of the OpenAI and by extension the DALL·E model. Basically, in terms of environmental impact, Azure has been working hard to achieve carbon neutrality. Azure has goals to be run by “100% renewable energy by 2025, water positive by 2030, and zero waste by 2030, and net-zero deforestation from new construction” (Microsoft, 2022). This information leads to the conclusion that by the extension of Azure’s goals of sustainability, the DALL·E model has also defaulted to valuing the environment. Overall, using Latour’s definition of prescription, the technology of DALL·E values open, easily accessed data as well as environmentally sustainable server infrastructure.

On the other hand, we can further define this technology by prescription by what or who it discriminates against. Specifically, the model, in a way discriminates against artists because it nonconsensual uses their artwork for training purposes. Did artists ever release their art on the internet to be analyzed by robots or did their consent extend only for human consumption? Currently, there is a lawsuit against Microsoft, GitHub, and OpenAI for “violating copyright

law” (Vincient, 2023) for developing code from open source resources. Similarly, AI art generators are being sued for “scraping” images from the internet “without consent of the original artists” (Vincient, 2023). Perhaps the outcomes of these cases will help define the extent in which AI is allowed to use data from the internet.

Latour next defines technology with the concept of circumscription. This part of technology affects users at certain times. Specifically, in this instance, users are affected when they have access to computers and the internet. In order to use the DALL·E model, one can easily go to OpenAI’s site and use a web-based API tool. This mostly limits the technology’s effect on people when they are on the internet to view or create AI generated images. However, this type of technology extends this circumscribed scope when some creators submit creations to art competitions (Roose, 2022). Also, people’s access to this technology is also dependent on their economic privilege. Those who have extra money to spend might benefit from this particular model’s API due to DALL·E’s access limits. After a user reaches a limit of free image generation (defined by the number of credits per user over a certain period of time), they must pay to generate more images in order to avoid waiting periods.

Finally, we can use Latour’s definition of description to help define the DALL·E model. The primary mechanisms of the DALL·E model were defined in the previous section, but to summarize, the model was trained to create images based on 12-billion input parameters and “250 million image-text pairs” (Ramesh et al., 2021). After training, the model is able to use the frequentist’s probabilistic theory to guess what image the user wants based on user text input. To use this technology, one registers an account with OpenAI and then can generate a certain amount of images per time period for free.

## Research Question and Methods

This research endeavors to answer the question: Do people's perception of AI generated art vary depending on their understanding of the model used to generate the art? This question is important because it could help identify specific issues people have with AI generated art. After specific issues are identified, software engineers can then approach a solution to develop better models. This particular study focuses its analysis with the DALL·E model.

The data was collected through surveys that were distributed to mostly students; although, other members of the general public were welcome to answer the survey. The survey depicted an image generated by the DALL·E model using the following text prompt: "Cyberpunk image of a guy smoking outside of a bar."



*Figure 1.* Prompt: Cyberpunk image of a guy smoking outside of a bar.

The criteria for analysis was a rating system (1-5) where the image was to be rated in categories described by Hong and Curran (2019): originality, composition, development of

personal style, experimentation or risk taking, aesthetic value, and successful communication of ideas. Participants were asked to rate the image before and after the explanation of the DALL·E model. Follow-up questions were also asked in an attempt to determine why participants changed their rating of the image if their ratings changed at all. Overall, the results of the experiment were analyzed using the Wilcoxon signed-rank test with Pratt's signed-rank zero procedure at a 0.05 rejection threshold. This test was used to reveal if there was a significant difference in rating AI generated art when subjects gained a better understanding of the DALL·E model's mechanisms. More qualitative follow up information was analyzed through close readings of subject responses to qualitative questions (see Appendix).

## **Results**

Perception of AI generated art did not change with a better understanding of the DALL·E model's mechanisms. Overall, 45 responses were collected, and 25 responses were used in the Wilcoxon signed-rank test. The responses selected for the test were from participants in the survey who claimed they either "read the summary of how DALL-E Works" or "clicked on the link(s) and reviewed information about the DALL-E model." Since the goal of this study is to understand if a person's views about AI art change after learning about a model, participants who explicitly state they did not do research about the model were excluded from this test. The rejection threshold for the test was 0.05 and p-values for the categories of originality, composition, development of personal style, experimentation or risk taking, aesthetic value, and successful communication of ideas were calculated to be approximately 0.8957, 0.04142, 0.7884, 0.7574, 0.08326, and 0.3944 respectively.



To start the test, the following null hypothesis was established: there will be no significant difference of ratings for the image before and after the model is explained to the participants. Once data for the experiment was collected the Wilcoxon signed-rank test was chosen because we could assume that data between subjects is independent and that each paired rating is from the same subject; hence variances are equal between the two sets of ratings (Ford, 2017; JMP Statistical Discovery). I could not prove that the data was normally distributed; therefore, a paired t-test was not appropriate (Ford, 2017; JMP Statistical Discovery). Because the data resulted in multiple zero values after a difference was computed between the paired data, Pratt's signed-rank zero procedure was utilized. This method allows for zeros to be included when ranking samples ("Wilcoxon signed-rank test," 2023). The rejection threshold for the test was predetermined to be  $p < 0.05$ . This specifically, means that if our p-values result in a value less than 0.05, then we are able to reject the null hypothesis and conclude that there is a difference in ratings for an image before and after the model is explained. However, since no p-values fall within the rejection threshold, we are able to fail to reject the null hypothesis and conclude that there is no significant difference between the ratings of the image. Thus, to answer the original research question: people's perception of AI generated art did not change with a better understanding of the DALL·E model's mechanisms.

However, although I cannot make conclusions that people's perception of AI generated art will change with a better understanding of the DALL·E model's mechanisms, there were interesting responses from some participants as to why they altered their ratings. Before discussing these responses, it is important to understand that these answers do not nullify the results from the Wilcoxon signed-rank test. The results from the Wilcoxon signed-rank test attempt to make conclusions about a population through sampling opinions. So, the results from

this test indicate that generalizations that deviate from the null cannot be made with the current sample.

Anyways, from the survey, one subject said that they lowered the ratings for originality and risk taking because the model is trained to “give you what you ask for and not go outside of the box.” The participant also claimed that since it is art from a computer, it should be of less value than art from people since it fails to capture “inspiration and creation.” Another participant agreed with the sentiment that AI is not able to have original thought or experiment since it is not sentient. Similarly, ratings for originality were lowered by another person because the art was not fully completed by a human. A subject dropped their ratings for aesthetic value and successful communication of ideas after remembering that they do not need to avoid hurting the feelings of the AI.

Some other comments from participants with regards to AI generated images relate to how the images were generated. Particularly, some participants felt like the AI was stealing data when generating images due to the lack of explicit permission from the artists. In the survey, one participant stated that “AI art seems like theft of other artists’ work.” This participant stated that they felt this way because AI is unable to use things such as “life experiences” and “imagination” for inspiration in addition to works from other artists. With these additional components of inspiration, humans are able to be inspired by others without “theft.” When another participant was reminded that an AI generated the art from the survey, they felt that “the originality of it [was] totally gone.” Others had similar attitudes where they felt that “art is something that expresses human emotion” and a computer or AI is unable to do this. Participants also had other negative sentiments about AI generated art such that the created art is “totally soulless” or an “atrocious.”

## Discussion

Relating the results of the study back to Latour, connections between some of the qualitative responses and the definitions of prescription and description can be made. Recall that Latour's definition of prescription was used to assess who the technology discriminated against. It was determined that DALL·E and other AI image generation models discriminate against artists. Particularly because artists did not give consent for their work to be used during model training. With regards to AI art generation, some have voiced that artwork is being stolen as it is freely and "quietly scraped from the web" to train "artificial intelligence models" (Enking, 2022). To Enking (2022), this is especially problematic because it could "take work opportunities away" from artists. Similar sentiments can be seen from participants in the results section as they describe DALL·E as "stealing" work from artists. Further, describing DALL·E with Latour's definition of description we see an inherent limitation of the technology: it is not solely human generated. Of course, in order to use DALL·E, a human must submit a prompt and run the algorithm; however, the actual creation of the image is executed by a computer. This characteristic of the technology is noted as a drawback by several subjects as they claim AI cannot generate original art that takes risks due to the computer's lack of self.

Other studies have also found that participants tend to devalue artwork when assessing art they think is AI generated. In one study, participants who had the bias that "AI cannot produce art" often rated art lower when they thought that a model had generated the image they were viewing (Hong & Curran, 2019). Overall, Hong and Curran's (2019) study concluded that AI and human generated art are distinguishable; however, there needs to be further research to evaluate what makes the works distinct so that AI artists are able to improve (Hong & Curran, 2019). Next, researchers Lima et. al. (2021) found that, when participants acknowledged AI's

“lack of mind,” it negatively influenced how they evaluated the artwork. Similarly, in this study, some of the participants found the art they viewed as less valuable when they remembered that the model/computer is not sentient and lacks feelings. Further, Lima et. al. (2021) considers the assessment that AI generated art may “lack the meaning necessary to be considered art” since “meaning can only emerge from human artistic communication.” Interestingly, Yamshchikov and Tikhonov (2018) both try to address the communication issue that AI generated art may have. They find that if a model can have an “unbounded decision space” with a “diverse set of outcomes,” then it can lead to a more “personalized experience” that humans may consider “interesting and insightful” (Yamshchikov & Tikhonov, 2018). Essentially, having models produce artwork with an “unbounded decision space” and a “diverse set of outcomes” may lead to AI generated art that is more human-like (Yamshchikov & Tikhonov, 2018). According to Hong and Curran’s (2019) study, if models are able to produce images that are more human-like, then people might value them more since participants, in their study, on average rated human art higher when they were unaware of who the artist was.

This study had a variety of limitations that should be considered when evaluating the results. Firstly, a larger sample size of the study would have been ideal for assessing the research question at hand. The sample size of  $N = 25$  does not diminish the outcomes from the Wilcoxon signed-rank test; however, more samples might be able to demonstrate the opinions of the population better. The study also only asked participants to consider one image generated from the DALL•E model. This is a limiting factor in the study as participants might have changed their ratings if they were given a different image. Only one image was given to participants so that they wouldn’t experience decision fatigue while answering the survey. A final limitation to consider is that the explanation of the DALL•E model was brief. It would have been ideal if a

more detailed explanation of the model was given; however, participants' lack of prior knowledge about general machine learning and artificial intelligence concepts might have made the more detailed explanation ineffective. In essence, since the study attempted to obtain opinions from people with a variety of backgrounds, the explanation of DALL·E was basic so that it can be understandable to many people.

If this study were to be conducted again, there are a few things that I would consider doing differently. Firstly, I would consider conducting a study using participants from a machine learning or artificial intelligence classroom. I would ask them to take the survey before and after their course term. However, I would use more sophisticated language to explain the DALL·E model in the second survey because a more detailed explanation of the model should make sense to more students after their coursework on the subject concludes. Also, I would create a focus group with some of those students to get a better qualitative understanding of their survey responses. The use of focus groups may create more accurate conclusions about the subjects' perceptions of AI generated art. However, when we conduct the study in this way, we have a biased set of participants. Since the set of participants is biased, results may not capture the opinions of the general population in regards to the research question.

General ideas from this research will be used to advance how I conduct myself as an engineer. Because, although I am uncertain if I will pursue a career in machine learning or artificial intelligence, broad concepts from this research can be utilized for advancing my engineering practice. Particularly, this research has influenced me to reconsider how I develop software in the future. For instance, when developing software, I can consider who the software discriminates against and how to limit this bias. Also, I can consider how the algorithmic

techniques a software utilizes might influence the values it holds. Then, I can evaluate the values and decide if they are values that should be altered.

## **Conclusion**

This research was able to recognize some key issues that people may have with AI generated artwork. The identification of issues is the first step in developing a solution to improve existing technology. The next step for others will be to pinpoint which attributes of AI generated art make the art distinguishable from that of humans. The identification of these attributes will allow models to improve their art generation techniques. Overall, improvements to AI generative art will have the potential to completely alter how the medium of art is viewed and valued.

## References

- Brockman, G. Sutskever, I., & Altman, S. (2016). Open AI and Microsoft. Retrieved from <https://openai.com/blog/openai-and-microsoft/>
- Enking, M. (2022). Is Popular A.I. Photo App Lensa Stealing From Artists? [Smithsonian Magazine]. Retrieved from <https://www.smithsonianmag.com/smart-news/is-popular-photo-app-lensas-ai-stealing-from-artists-180981281/>
- Ford, C. (2017). The Wilcoxon Rank Sum Test. Retrieved from <https://data.library.virginia.edu/the-wilcoxon-rank-sum-test/>
- Hong, J. W., & Curran, N. M. (2019). Artificial Intelligence, Artists, and Art: Attitudes Toward Artwork Produced by Humans vs. Artificial Intelligence. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 15(2s), 1-16. <https://doi.org/10.1145/3326337>
- JMP Statistical Discovery. The Paired t-Test. Retrieved from [https://www.jmp.com/en\\_us/statistics-knowledge-portal/t-test/paired-t-test.html#:~:text=Paired%20t%20test%20assumptions&text=Subjects%20must%20be%20independent.,be%20from%20the%20same%20person.](https://www.jmp.com/en_us/statistics-knowledge-portal/t-test/paired-t-test.html#:~:text=Paired%20t%20test%20assumptions&text=Subjects%20must%20be%20independent.,be%20from%20the%20same%20person.)
- Lai, C., Ahmed, S., Dubinsky, D. & Maver, C. (2022). AI is harming our planet: AI's staggering energy cost. Retrieved from ["https://www.numenta.com/blog/2022/05/24/ai-is-harming-our-planet/#:~:text=How%20much%20energy%20does%20AI,to%20only%20one%20training%20run."](https://www.numenta.com/blog/2022/05/24/ai-is-harming-our-planet/#:~:text=How%20much%20energy%20does%20AI,to%20only%20one%20training%20run.)
- Latour, B. (1992). Where are the Missing Masses? The Sociology of a Few Mundane Artifacts. In W. E. Bijker, & J. Law (Eds.), *Shaping Technology/Building Society: Studies in Sociotechnical Change* (pp. 225-258). Cambridge, Massachusetts: MIT Press.
- Lima, G., Zhunis, A., Manovich, L., & Cha, M. (2021). On the social-relational moral standing of ai: an empirical study using ai-generated art. *Frontiers in Robotics and AI*, 8. doi:10.3389/frobt.2021.719944
- Mazzone, M., & Elgammal, A. (2019). Art, creativity, and the potential of artificial intelligence. *Arts*, 8(1), 26. doi:10.3390/arts8010026
- Microsoft (2022). Azure Sustainability. Available at: <https://azure.microsoft.com/en-us/explore/global-infrastructure/sustainability/#overview>
- Salkowitz, R. (2022). AI Is Coming For Commercial Art Jobs. Can It Be Stopped? [Forbes]. Retrieved from

<https://www.forbes.com/sites/robsalkowitz/2022/09/16/ai-is-coming-for-commercial-art-jobs-can-it-be-stopped/?sh=11c9baf454b0>

Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., . . . Sutskever, I. (2021). Zero-shot text-to-image generation. *ArXiv: Artificial Intelligence*.  
<https://doi.org/10.48550/arXiv.2102.12092>

Roose, K. (2022). AI-Generated Art Won a Prize. Artists Aren't Happy. [The New York Times]. Retrieved from  
<https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>

Vincient, J. (2023). AI art tools Stable Diffusion and Midjourney targeted with copyright lawsuit. Retrieved from  
<https://www.theverge.com/2023/1/16/23557098/generative-ai-art-copyright-legal-lawsuit-stable-diffusion-midjourney-deviantart>

Wilcoxon signed-rank test. (2023, January 30). In *Wikipedia*  
[https://en.wikipedia.org/wiki/Wilcoxon\\_signed-rank\\_test#:~:text=The%20Wilcoxon%20signed%2Drank%20test,%2Dsample%20Student's%20t%2Dtest.](https://en.wikipedia.org/wiki/Wilcoxon_signed-rank_test#:~:text=The%20Wilcoxon%20signed%2Drank%20test,%2Dsample%20Student's%20t%2Dtest.)

Yamshchikov, I. P., & Tikhonov, A. (2018). I feel you: What makes algorithmic experience personal? *Electronic Workshops in Computing*. doi:10.14236/ewic/evac18.32



## Appendix

### Survey Questions

1. Did you read any information from this survey? (Select all that apply)
  - a. I read the summary of how DALL-E Works
  - b. I clicked on the link(s) and reviewed information about the DALL-E model.
  - c. I did not read the summary of the DALL-E model nor did I click the link(s) to review information about DALL-E
2. How much did you know about the DALL-E model/AI generated images before doing this survey? (Select all that apply)
  - a. I knew nothing about DALL-E
  - b. I've heard about DALL-E
  - c. I knew some things about DALL-E
  - d. I knew a lot about DALL-E before this survey
  - e. I've heard about AI generated images before
  - f. I've never heard about AI generated images before
3. If you changed the ratings of the image, which ratings did you change and why did you change your rating?
4. Do you have any opinions about AI generating art? If yes, could you summarize them?
5. How much prior knowledge about DALL-E or AI generated images did you have before? What did you know prior to the survey?
6. Did you do any external research about DALL-E or AI generated images during the survey? If yes, would you be willing to share the resources?
7. Do you find anything particularly intriguing or unique about the image?
8. Does the image feel familiar?
9. Is there anything you particularly liked or disliked about the image?