Creating Art Using Generative Adversarial Networks

Henry Abbott ‘21

A.I. for the Humanities

Introduction

I set out to explore how artificial intelligence generated artwork compares to human created artwork. Using CycleGANs, my goal was to create art, and then conduct a Turing Test in order to determine whether or not one could distinguish between A.I. and human art. Using this test, as well as my own observations, I planned to analyze where CycleGANs succeeds in creating art, as well as where it struggles. Following this analysis, I thought more about what the results mean for art moving forward, and if artificial intelligence can truly create art.

About the Technology

In order to create art using Artificial Intelligence, I turned to CycleGANs, a type of generative adversarial network (GAN). Using CycleGANs, we can input two datasets of images, and after training the model, which consists of four neural networks, it is able to generate a brand new image. Two of the neural networks are generators that generate new images similar to the images in their associated dataset. In the training process, one generator uses an image from the other dataset to create a new image, and then the other generator uses that new image to generate, or reconstruct, the original image. This is also done in the other direction. The generators are trained on these image reconstructions by calculating the loss between the image and its reconstruction, and then through gradient descent, using backpropagation to calculate the gradient, the CycleGANs updates its weights. Once the weights are updated from this process, the other two neural networks, called the discriminators, are trained. An image is put through its dataset’s discriminator and the loss is calculated. Then, a generated fake image is put through the discriminator, and this loss is averaged with the loss from the real image. Using this loss, the weights of the discriminator are updated, and this whole process is done with both discriminators. Once all four of these neural networks are trained, we can use CycleGANs to generate brand new images. This is a sophisticated network, and after working with it and debugging for a long time, unfortunately it was a struggle to successfully run the program, and we were unable to generate images. In light of this problem, a StyleGAN was the next best option to create computer generated art. While this GAN was not as complex, along with possibly being less effective, it still was able to generate some very interesting art from a large dataset of images.

The Turing Test

After creating images using the StyleGANs, I put them to the test using a Turing Test to determine how they stacked up against human created art. For each image, I used Yandex to search for art similar to the image, and after choosing the two closest options, I showed individuals the three pictures and asked them to tell me which one they thought the A.I. created. I did this with 28 Kenyon Students of all ages and backgrounds, having each person go through all four groups of images individually in order for them to not be influenced by what other participants thought. I have the computer generated pieces on the left for all four groupings, but when conducting the Turing Test, the images were randomly ordered.

StyleGANs Generated Images and Comparisons

StyleGANs, 46.4%

Rezanova-Velichkina Olga, 39.3%

Malda Muizule, 14.3%

StyleGANs, 67.8%

Ernest Lawson, 21.4%

Grigorij Musatov, 10.8%

StyleGANs, 71.4%

James Green, 21.4%

James Whistler, 7.2%

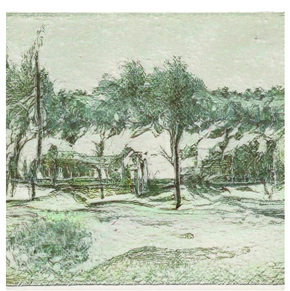

StyleGANs, 78.6%

Edward Mitchell Bannister, 7.1%

James Ensor, 14.3%

Each row above is a grouping of 3 images that was shown to each participant. Under each image, I have listed the artist of the piece, as well as the the percentage of the participants that thought the piece was the A.I. generated image of the three. The StyleGANs images had the highest percentage in all groupings, though this percentage varied greatly from one grouping to the next.

Analysis

Judging by the results of the Turing Test, StyleGANs has not passed. Aside from the first grouping of flower paintings, where the artificial intelligence fared very well, most participants were able to distinguish between the A.I. and human pieces. Many participants picked out oddities in the StyleGANs creations, and were then able to correctly pick out that image. These oddities included the unusual ear of the man which led to a 71.4% success rate for the participants, or the A.I.’s struggle to depict structures like houses in the bottom image, which led to a success rate of 78.6% for the participants in that grouping. Some were fooled by these imperfections, viewing them instead as abstractness or style. However, most viewed them as giveaways, which led to a very high success rate for participants. StyleGANs was very successful in the top grouping, though, as it was able to depict flowers in a very convincing fashion. It is clear that this technology succeeds in generating more vague, random images like flowers or blurry landscapes, but struggles when generating detailed objects such as people or buildings.

Conclusion

While StyleGANs struggled to pass the Turing Test, there are still some very interesting observations to be made. After generating the images, I searched for similar images, and for some of them, I found it very difficult to find images with a similar style. Some may say this is because A.I. can’t produce art like humans can, but I would say this is because the style of the A.I.'s art is different from conventional styles. I was unable to find pictures for the bottom grouping that were of the same style and coloring, possibly the reason why participants were so successful in picking correctly. Not only can this technology be used to generate art, but as we can see, it can be used to think of new styles that may not be as common. Moving forward, some may be hesitant to say that a computer can create art, but even a relatively simple StyleGANs can and has. This is certainly not even the most sophisticated or successful technology right now, and it is extremely reasonable to think that in the very near future, we will begin to see far more advanced generative networks. While I do not think that artists will ever be replaced with technology, as I do believe that intent, emotion, and expressiveness are important to both artists and admirers of art, I see no reason to think that A.I. art won’t continue to advance and grow. It can emulate some styles better than others, but moving forward, it will improve quickly to the point where it will be capable of creating art of any style. It can also create new styles we haven't seen before. Through my work with GANs, I have seen first hand what artificial intelligence can generate, so while art will not be solely created by technology in the future, it is undeniable that technology has a future in art.

Acknowledgements

Thank you to Professor Jon Chun and Professor Katherine Elkins for assisting me in this exploration of artificial intelligence generated art. Both Professors were vital in the completion of this project, especially in the execution of CycleGANs and StyleGANs.