ART AND MACHINE LEARNING
CMU 2022 SPRING
PROJECT 2

The Poetic Interpretations of a Computer

Group 21: Sana Gill

DESCRIPTION

This project is centered around the visualization of poems. The idea being that a user could enter the text from a poem and the output would be an image. The large text of the poem is first processed and then fed into vqgan, stanza by stanza, where each line is added as a separate descriptor. The final image, created from individually visualizing and combining each stanza, should then be representative of the entire poem. Poetry often consists of hallucinatory imagery, metaphors that when interpreted literally may be illogical, and in addition disparate elements from stanza to stanza. These very qualities make it ideal to generate novel, eye-catching surrealist artworks.

Concept

When you read a poem you often create images in your head as you read. Having a computer generate visualizations provides a new interesting lens at which to look and interpret the poem. The visualization may be similar to what you were seeing as you processed the text internally, or it could be wildly different. Transforming one art into another artform creates a new piece entirely that can be enjoyed in conjunction or even separately. Initially I was inspired to create computer generated surrealist works. As I went through the process though I struggled to create images that captured my intentions and that had both an intriguing style and substance. Another idea I had prior to starting was perhaps doing text to image generation. I decided that I could possibly combine the two ideas, and that is when my project slowly transformed from painting generation, to text to painting, and finally a focus just on poems, which seemed to capture the style that I had been intending for all along. Surrealist pieces often have such a response from their audience because the composition, contrasting images, and distortion of reality subverts our expectations. This is why I found poetic text to be so well aligned with surrealist paintings since oftentimes the text may have an overarching resonating statement that is rooted in reality, but the sentences received individually can have that same dreamlike quality and distance from literal interpretation. The results ultimately are an insightful look at a computer's visualization of poetry and the images themselves are standalone surrealist pieces.

Technique

In order to have the final output image capture the poem in its entirety I had the poem separated by both stanza and line. When a user inputs a poem, it is separated by stanza so that each stanza can be run on vqgan separately. When a stanza is inputted as text for vqgan, it is written with each line separated by a '|'. This allows for the lines to be individually interpreted and blended together in the output for that stanza. When a stanza results in the output of an image that image is fed is added as a base image for vqgan so as to influence the output of the next stanza. This is repeated again and again, an image being created for a stanza used as influence for the next stanza and so on. The final image, created from individually visualizing and combining each stanza, should then be representative of the entire poem.

Process

- First I attempted to use Style Transfer by Relaxed Optimal Transport and Self-Similarity (STROTSS) Style Transfer. This would have two images inserted, where one is a random image of my choosing and the other was a surrealist painting whose style I wished to be imposed upon the previous image.
    - The resulting image was oftentimes clear, intriguing, and very obviously of the style of the surrealist painting that was provided. The downside I found is that every image would be very similar to each other for a given surrealist painting used because that one style is incorporated. This was not exactly what I was looking for. I wanted for a random image to receive a random surrealist painting equivalent of that image, and not for example always getting reimagings of Salvador Dali's Persistence of Memory painting (that was one such image used as the 'style image')
- I then tried using v-diffusion-pytorch. Instead of two images, a text prompt is given and the resulting image interpretation is outputted.
- After some tinkering with the above choice and changing the text prompt, the number of diffusion timesteps, the initial timestep, etc. the resulting images looked exactly as I had been intending at the start of the project.
    - Though I liked the images I would get I wasn't satisfied with having each image not being an interpretation by the computer, but simply an artistic idea that I had, inputted, and formatted so it would be visualized in some way that looked vaguely as I had intended.
- This inspired me to incorporate an idea I had earlier about poems to images, so I started looking for ways to use the poem text as prompts for the image output.
- I then tried using VQGAN + Clip and also CycleGan. Ultimately after using both I found the outputs of VQGAN + Clip to be more what I wanted because the outputs, though not always reliable had more of a hallucinatory and dream like quality to them that I was looking for
- I then worked to modify the code by having the poems inputted and processed in such a way that they were entirely incorporated. When I first used the whole poem as a text input the output was not entirely encompassing and focussed instead on one random element.
- In the next iteration I had the text being inputted as sentences separated by '|', but the approach did not exactly represent the poem either because the elements would be all there but not incorporated as they had been interacting in the poem.
- I then tried separating by stanza since the stanzas are representing clear delineations in the poem. To incorporate the separate stanza images together I decided to use the following image as the influencing image for the next stanza.
- I then did several iterations of trying and using different weights for how much influence the previous image should have. I stopped at the given weight once I was satisfied with the outputs.

Reflection

I chose the final result because I loved how the visualization turned out. The interpretation seems to follow the poem, but also is imaginative and with a startling composition. It draws your eyes with the distortions and incongruence of images, colors, and textures. I was really satisfied with how the idea for

this project translated into the final work. I think the technique could always be further tweaked to refine the process of translating the poem and perhaps make it quicker and with less extraneous images. For the time and scope of this project I think it turned out really nicely and the paintings themselves are interesting computer generated surrealist paintings. Initially in the process I had focused on simply creating the paintings, but as I continued to work it evolved into the transformation of text into the paintings, and then later solely on the transformation of poems into paintings. As I worked through the process of creating the art the exact idea I had got more and more specific, allowing me to have a clearer focus on what my end goal was. Once I had the idea fully formed, the trial and error process was mostly just tinkering around with what would be the best algorithm to use and how I could adjust the parameters in order to refine the image. That process did not end up taking as much time as I had anticipated. Overall, I learned quite a bit on how to utilize GCP and Collab effectively. Prior to this project I had used collab but not to this degree and running projects that needed extensive training and could take an exorbitant amount of time if not for utilizing the gcp vm. I also learned a lot about the different tools that I could use to generate images via text and images via other images. From working on the project I saw the advantages and disadvantages of using the different tools that we had covered in class more clearly.

RESULT

The result is a series of images, where each image is the result of inputting a specific poem. With the code any given poem can be visualized as an image and the chosen images are just some specific examples of images that were generated from specific poems I inputted. Created images can be found at the following link:
https://drive.google.com/drive/folders/1CWyx12xcBunbfe8w_-KkssIgmYVbygDi?usp=sharing

CODE

https://colab.research.google.com/drive/17N9qONl4CxmzBxYKETdb-BVS8tXJ1pOy

REFERENCE

N/A