ART AND MACHINE LEARNING
CMU 2019 SPRING
PROJECT 1





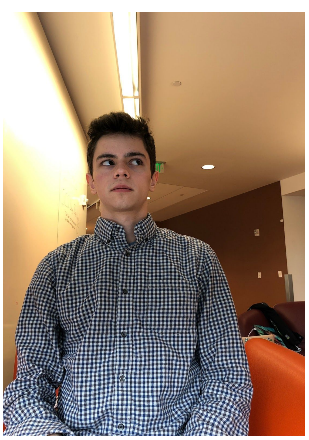


Trio_Style_Tranfer

Ananya Chandra, Edward Goebel, Vera Schulz

Our Concept:

Imagination:

All three of us had different concepts of what art meant, and each of us had a different interpertation of it. We wanted to identify artists that aligned with parts of our personal identity. These artists would then, have a connection with us through their techniques and artistic representations. We wanted to marry our personal identity with art, to show how art taste can represent a deeper personal bond and experience in our globalized world where identity is always shifting and being redefined. Some questions we asked were: If style transfering a specific artists style would make it less unique and more prone to becoming a homofally? How different are the facial reactions of people when looking at a specific artwork? (Why?)

Each of us ascribe a specific personal identity/connection to the artist.

- **Vera**, specifically chose *Wassily Kandinsky* because she likened the chaos of his strokes and the vibrance of his colors, to her creative process and inquiry. She also appreciated the disarray and undefined aspects of his painting, believing that this highly variable painting could beinterpreted by anyone.

- **Ananya** chose the style of *Jamini Roy* because of the importance of this artist's style for Ananya's identity. Having grown up with Jamini Roy style artwork on her walls at home, the artwork represents her background and culture.

- **Eddie** chose the style of *Frank Stella* as he was especially captivated by the artist's use of mesh-like geometry used throughout the selected work, *Juam*. To Eddie, the style represents a busy, hectic interaction of overlapping abstract objects. Looking closer, elements become more familiar and can have an identity assigned to them while remaining cluttered, scattered, and ultimately unformed. The work reminded him of his own perception of the workings of his mind.

Meaning:

The meaning we are looking to communicate with our work is the power of identity. By taking paintings that represent us in one way or the other, and distorting them using style transfer, shows the power of redefining something that seems so defined. Then, after the painting is distorted, it is masked on our faces, showing how easy it is for identity to shift and completely change, but in some ways it stays the same. In our final work, we display our three distinct faces, morphed with our chosen style, as the work flickers among phases of style transfer. Despite each of us having our own faces, and our own identities, it is with our unique individuality that we contribute to the whole. Our identities do not tear us apart; rather, by expressing and sharing our identities together we can create something beautiful. In doing so, machine learning has helped us use our uniqueness to bring us closer together in an expression of our identities.

Technique:

In order to accomplish our goals, we set out to style transfer our respective representative styles onto our faces and view the progression of the style transfer at each of the iteration, combining them into a moving picture (GIF). As the model progressed on our photo, we wanted to view the emergence of the style into our respective identities.

Before diving into the project, we debated what we wanted to capture, and what form the data would take that would be fed into the machine learning algorithm. We decided we wanted to capture the our identities and then merge them with an artists style to create a composite. The data would be fed into the algorithm by pictures. Since the styles are representative of our backgrounds and cultures, we want them to be intensely displayed in our personal photos. In order to improve the intensity, we wanted to experiment with the style transfer intensity parameter and content intensity parameter. The goal was to decrease the intensity of the content and increase the intensity of the style.

Additionally, we wanted the observe how many iterations it took for the final output to “converge” visually speaking. With our experiments, we kept the iteration number to the number that reflected the convergence of the algorithm. We wanted to see how stark the contrast was between iterations and then wanted to chose the pictures which has the most interesting results (in terms of style intensity).

Finally, we wanted to experiment with the content loss function using two different distance metrics: Euclidean distance and Manhattan Distance. With Euclidean distance, we used **(a-b)^2** where a and b represent the original content and the content we need to measure. With Manhattan distance, we utilized **|a - b|**.

We also looked into removing the final feature layer **(feature layer: block5_layer1).**

Finally, we looked at how the detailing of the style transfer was effected based off of the cropping of our images.

Final parameters we settled on for our style transfers:

```
iterations =10 #we will run 20 iterations during the optiomization. each iteration you get better value. this one is fo

# these are the weights of the different loss components
total_variation_weight = 0.8
style_weight = 2.0 #the best number authors got from their trials
content_weight = 0.015 #the best number authors got from their trials
```

Our tested content loss function:

```
# an auxiliary loss function
# designed to maintain the "content" of the
# base image in the generated image

def content_loss(base, combination):
    distance = K.abs(combination - base)
    return K.sum(distance)
```

Process:

We engaged in various trials when completing this project.

- **Trial 1**: Weight initialization
- **Trial 2**: Style Transfer intensity

Weight Initialization

One of the first trials we attempted was changing the weight initialization from the imagenet's pretrained weights to a random initialization.

```
# build the VGG19 network with our 3 images as input
# the model will be loaded with pre-trained ImageNet weights

model = vgg19.VGG19(input_tensor=input_tensor,weights='None', include_top=False)
print('Model loaded.')
```

However, that resulted in pretty much a grey mess (see below). This experiment gave us a clear understanding of the importance of using the learned Imagenet weights, and what we did not want our pictures to look like.

Iteration 0 **Iteration 2** **Iteration 4** **Iteration 6** **Iteration 8**

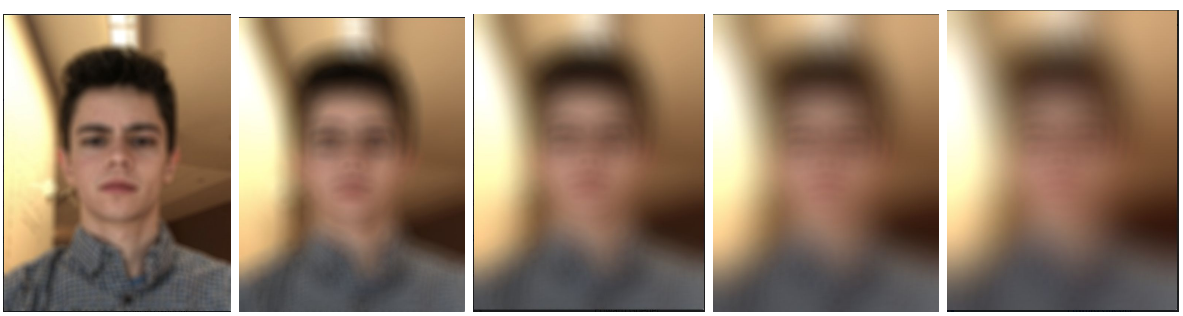

Crop and Style

We also resized our images and experimented with the effect of the resizing on the style that is applied to the image. Our goal was to capture more details in the facial features and how the style is applied to the facial features.

Below are our results:

In order to capture the detail of the style transfer onto the face, we decided to go with the second cropped version of the image.

Style Transfer Intensity

Style Intensity: 1.0

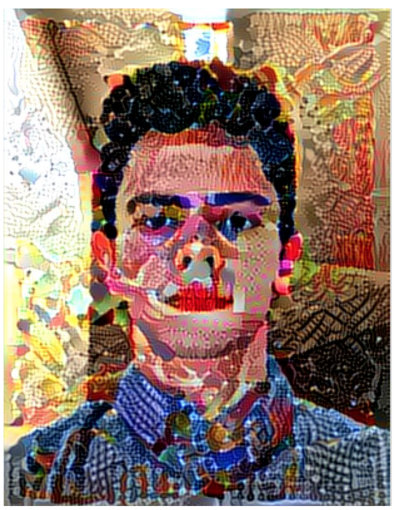

Style Intensity: 1.3

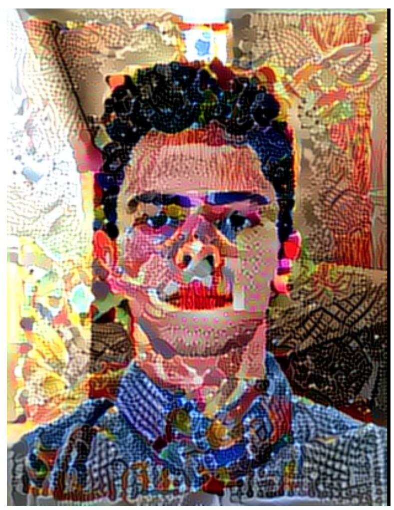

Style Intensity: 2.0

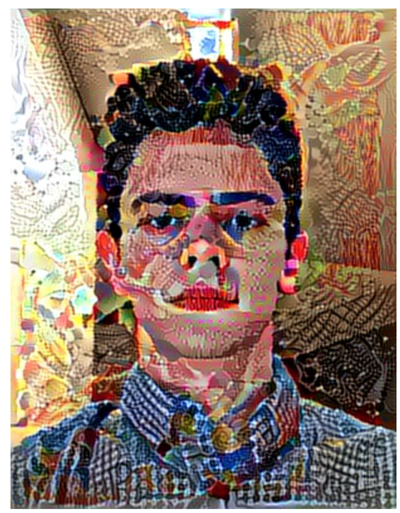

We chose to use the style intensity of 2.0 as it represented the intensity of our background and culture into our modern day identity.

Number of Iterations

We found that the iterations starting converging around the tenth iteration and decided to mark that as the parameter we would stop iteration afterwards.

Iteration #9 **Iteration #11** **Iteration #18**

<u>Content Loss Function</u>

We experimented with Euclidean distance and Manhattan distance formulas for our content loss function to see how it would affect the style transfer. The euclidean distance formula seems to bring on a bit of the color schematic of the style while the Manhattan distance formula maintains the RGB values and general tonality of the original picture. Because we wanted the exaggeration of the style, we opted to use the Euclidean distance formula.

Manhattan **Euclidean**

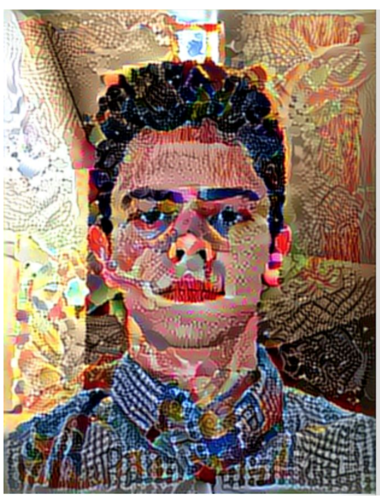

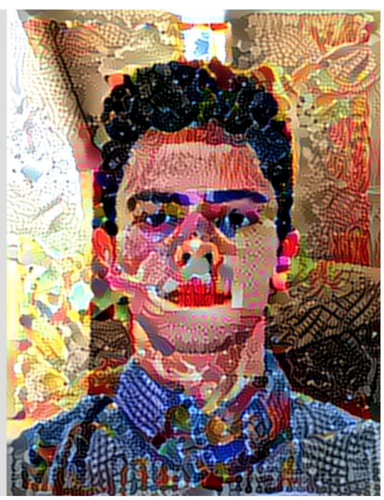

When we fed the pictures into the algorithm we realized that inputting them as PNG or JPG files did not matter, because the results were the same.

Giphication:

After we had the desireable results from the style transfer algorithim, we choose 3 pictures from the total images, being **1** (The first, having no style), **5** (the fifth, having some style) and **11** (maximum style).

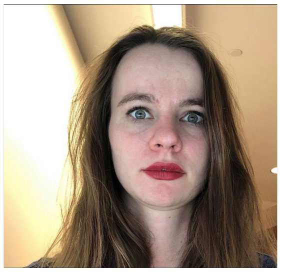

(1)

(5)

(11)

The *three* chosen images were then fed into on online gif maker: https://ezgif.com/maker & spliced together to make a gif (below) that demonstrated the changes in style.

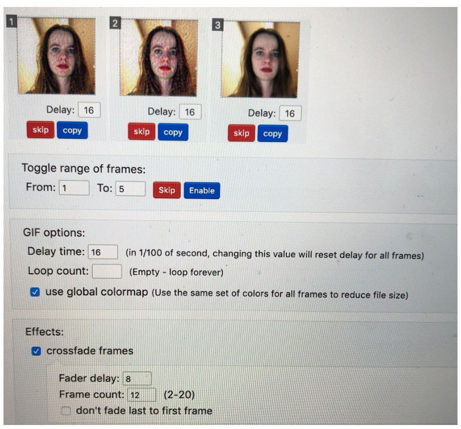

Setting the Loop/Fader/Frame Count

Giph

It was important to set the frame count at a fast enough speed that it could play through all 9 images but also slow enough that the viewer could see the style transfer happening with each individual person. The test was successful and then repeated with the 9 pictures to create the final composite gif.

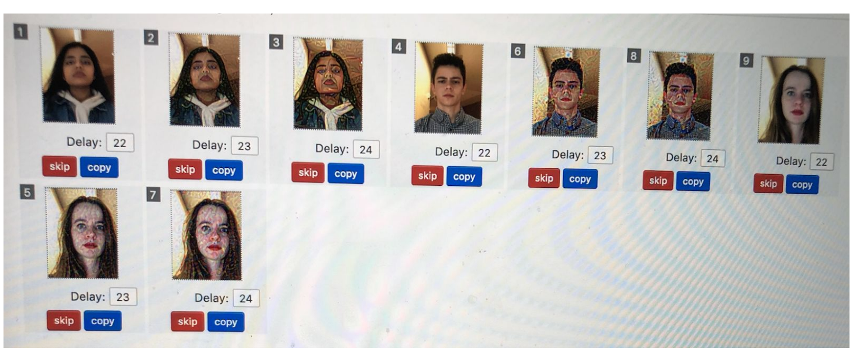

Result: Composite Gif of Ananya, Vera, Eddie

Final output to see the live GIFS
https://github.com/ananyachandra/StyleTransferAssignment1/blob/master/FinalOutput.mp4

Also the *three individual gifs* of everyone's face:

Vera

Eddie

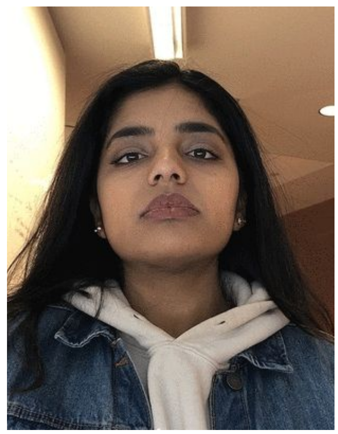

Ananya

Reflection:

We chose to only represent 9 images in our final gif, so that it would not be an extremely long gif and that it could be easily played back. We thought it was super cool that we were able to create an actual gif that showed the different style transfers we chose. We would have liked to experiment with even more intensity in the style transfer, and maybe choose a specific part of the artwork to apply, and see if the style would look more similar to the artists we chose.

Any new ideas that came into your mind, any reflections about trials and errors, and/or what you have learned etc?

- “It would be cool if the individual gifs looped better, as in we see it slowly style transfer, and then we see it "un-style transfer" and then the gif repeats so it’s always a fluid transition.” We implemented this suggestion (see individual gifs)

Future Steps:

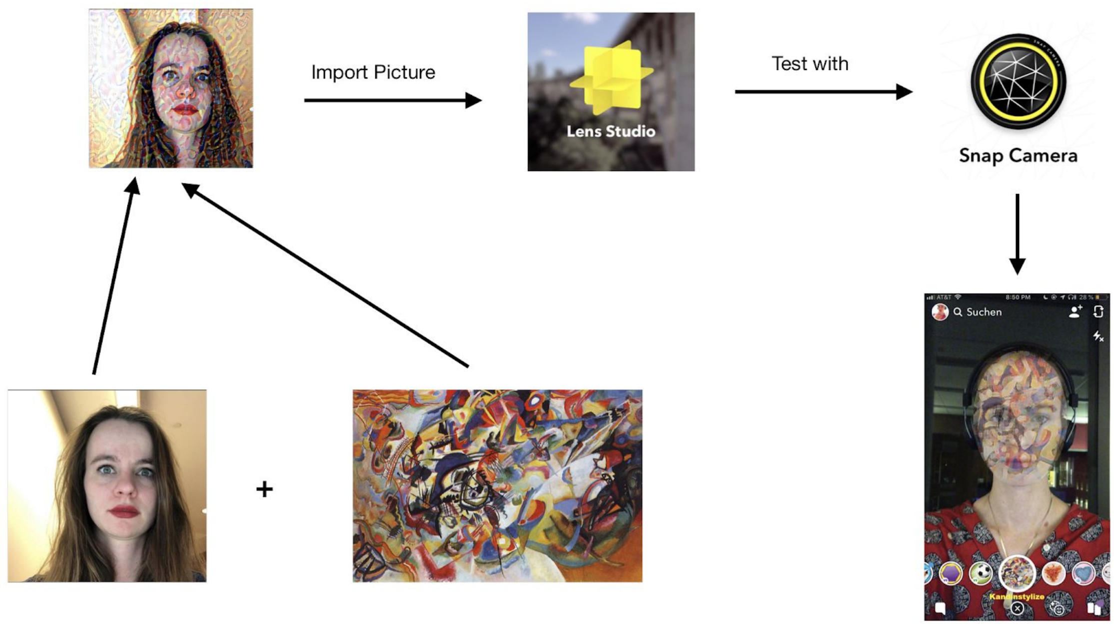

We would like to experiment with snapchat’s API and see if we can take one of our style transfer images or even our gif and turn it into an animated filter. Snapchat’s lens studio and snapchat camera would let us import the image/video and then create our own filter, which we could upload on the snapchat server (potential flow: Snap API > Image > Filter > Output).

Recognize Face > Style Transfered Image > Associated Emotion (chosen) > Emotion x style transfer

Ideally we would be able to have a snapchat filter that uses style transfer to alter the content style and weight depending on the facial expressions the person is making at the snapchat camera. This means that if the person changes their expression from sad to happy, the style transfer would change in style and weight parameters, correlating to the changed emotion.

I should be pretty evident to see the difference between a style transfer filter overlayed on a sad person's face vs a happy person's face. This functionality would also show the slight changes from a person who is very happy to content **vs** to extremely sad vs upset. The data that would need to be collected is all the different emotions a person can feel, and also the microexpressions. It would be important to initially focus on one specific person and see how the styles would shift.

Sources:

Trained Neural Networks and Snapchat

- https://community.wolfram.com/groups/-/m/t/1313034

Snapchat and Deepstyle:

- https://www.reddit.com/r/deepdream/comments/5s2y0z/snapchat_and_deepstyles/
  https://imgur.com/a/Q6nOq

References:

CODE (Github): https://github.com/ananyachandra/StyleTransferAssignment1

Technical:

[1] (Makow, Noah, Stanford University, Title: Exploring Style Transfer: Extensions to Neural Style Transfer)
http://cs231n.stanford.edu/reports/2017/pdfs/428.pdf

[2] (Gatys, Leon A, Neural Algorithm of Artistic Style
https://arxiv.org/pdf/1508.06576v2.pdf

[3] Kang Eun Sun - Github ArtML/Style_Transfer_Tensorflow
https://github.com/kangeunsu/ArtML/tree/master/Style_Transfer_Tensorflow

[4] (Singh, Manjeet, Medium Post, Artistic Style Transfer with Convolutional Neural Network, Sep 4, 2017)
https://www.cv-foundation.org/openaccess/content_cvpr_2016/papers/Gatys_Image_Style_Transfer_CVPR_2016_paper.pdf

Art

[1] Frank Stella, “Juam”
https://www.thedailybeast.com/larry-nassar-judge-aquilina-sentencing-victim-death-warrant?ref=scroll
https://www.tate.org.uk/art/artworks/stella-juam-p12327

[2] Wassily Kandinsky, “Composition VII”
https://www.wassilykandinsky.net/work-36.php

[3] Jamini, Roy, Various Artworks
https://www.culturalindia.net/indian-art/painters/jamini-roy.html