ART AND MACHINE LEARNING
CMU 2019 SPRING
PROJECT 2

Miscegenation

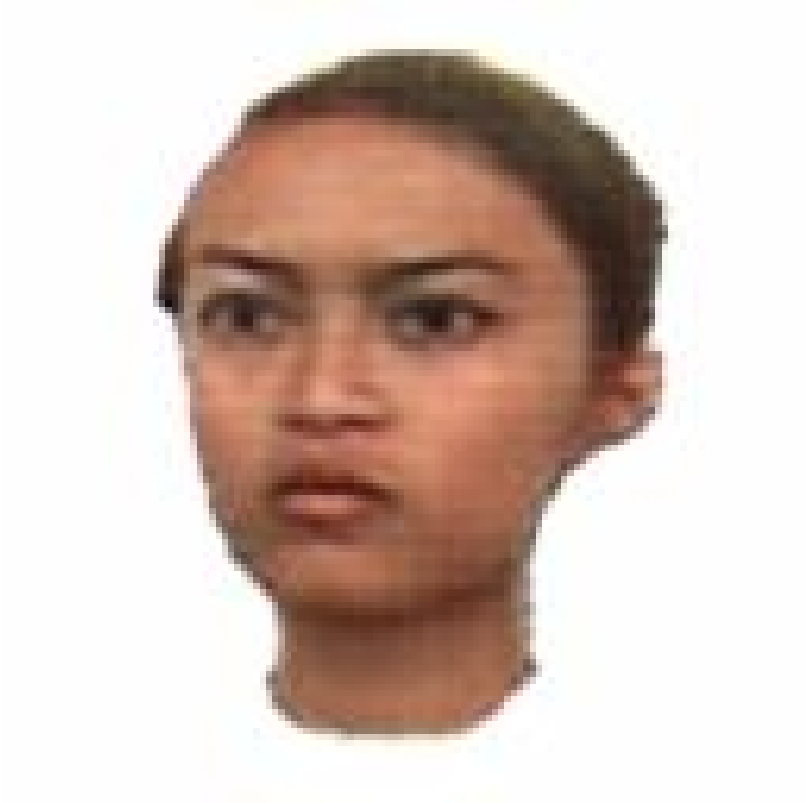

Stella Kim (hyunjunk)
Kevin Song (kmsong)
GitHub:
https://github.com/stellahyunjungkim/ArtML-project2?fbclid=IwAR0-aT9Lti8071vkrXfcU8A1LWb0Vg4XuummVG92BBMYD2vbPwiW0uEZ3EM

Introduction

Miscegenation noun : a mixture of races

The goal of our project was to explore different computer generated images for what human faces could look like in the future. In 1925, the famous philosopher Richard von Coudenhove-Kalgeri wrote, “The man of the future will be of mixed race. Today’s races and classes will gradually disappear to the owing to the vanishing of space, time and prejudice.”

Living in a tumultuous time of race and prejudice, the idea of one future mixed race seemed like an interesting topic to explore. Furthermore, following the past trends of immigration and increasing rates of interracial marriages, we came to the conclusion that the proportion of mixed-race beings would outnumber the proportion of “single-racial” beings, thus increasing our interest and the importance of the topic of miscegenation.

Approach

Our group utilized CycleGAN algorithms to obtain our results. The database of faces we used was from Tarrlab at Brown University, which contains multiple images for over 200 individuals of many different races (examples shown in figure 1). The three races we chose to “mix” for this project were Asian, African-American, and Caucasian. The approach we used for “racial mixing” was applying CycleGANs with unpaired datasets X and Y to be Asian images African-American images, respectively. We then took the resulting images, Z, and created a new unpaired database of Alpha and Beta to be Z and Caucasian images, respectively.

Figure 1. Example of Asian, African-American, and Caucasian Face (right to left)

Since we had to run the code twice, which takes about 18 minutes per iteration, we decided to shrink the iteration from 150 to 10. That reduced the training time to four hours. We might have gotten better and more clear results with higher number of iteration, but we figured out the 10 iterations were enough for what we have expected.
We wanted to save the resulting image, hence used batch size of 1 so that we got a single image at a time instead of a chunk of multiple images. This increased the training time, but simplified the whole process for us.

Expected Results

The expected results were more or less self-explanatory, which consisted of having some visible/identifiable combination of features from the three races. From a human's perspective, we expected the GAN to pick up on features such as skin-tone, face shape (round/long/etc.) and facial feature size/shape (eyes/nose/mouth/hair/etc.). Interestingly, while looking online, we found an article by Forbes depicting an artist's vision of what humans might look like in 100,000 years. It was interesting to note that this image only depicted obviously Caucasian faces, possibly ignoring the potential for mixed races. The faces were also near-comical in their enlarged features.

Figure 2. Designer Lamm's depiction of how the human face might look in 100,000 years

Our Results

Phase 1:

As described in our approach, our team began by using Asian faces and African-American faces as our initial unpaired datasets. For most of the images, it seemed as though a majority of the resulting faces consisted of Asian faces that had darker skin complexions (examples shown in figure 3). This was in coherence with our initial assumption that skin-tone would be a feature that the cycleGAN algorithm would pick up on.

Figure 3: Results from phase 1 ("mixing" Asian and African-American Faces)

Phase 2:

Once the resulting faces were saved from phase 1, we used this newly created dataset along with the dataset of Caucasian faces to produce our final results (examples shown in Figure 4). Generally speaking, we found that our final results produced slightly lighter-colored faces and hair, which again aligned with our initial expectations, while also sharpening some of the facial features such as nose and mouth. The shape of the face (i.e. cheekbones, jawline) also seemed to narrow more in our final results.

Figure 4: Final Results ("mixing" Asian, African-American, and Caucasian Faces)

Reflection

In our expected results, our ultimate hope was to be able to have resulting images of faces that could clearly identify as their own, unique race. Upon seeing our results, we found that each face would still be able to be potentially categorized as a specific, known race (i.e. Asian, African-American, Caucasian), despite each image exhibiting characteristics from all three races. This could partially be attributed to our dataset, which only had approximately 1000-3000 image files per race. Additionally, within each race dataset, there were duplicate images of individuals with varying lighting, perspectives, emotions, and disguises (e.g. glasses), which further reduced the total number of unique faces for each race. This in turn most likely reduced the fidelity of the trained discriminators in our cycleGAN algorithm.

References

Figure 1 URL :
http://wiki.cnbc.cmu.edu/Face_Place?fbclid=IwAR3GjQiSOlitTelYc9Dmfikh_kNCzOlkJbJ1F9dOONs9BrUq7oUXgDjKaEE
Figure 2 URL:
https://www.forbes.com/sites/parmyolson/2013/06/07/how-the-human-face-might-look-in-100000-years/#5cd8a4685a96