Project Final Report

Motivation/Overview

My project tackles the problem of finding a hairstyle that fits your facial features. Many people, including myself, have struggled to find hairstyles that work. A face has certain features that can point to a good hairstyle, and this project hopes to find those features and generate a completely new hairstyle for a given face. An individual's face takes one of eight identifiable forms(heart, square, pear, rectangle, round, oval, diamond and oblong)[1]. To get the full potential out of a haircut, one must take into consideration their specific face shape. Certain face shapes work with many hairstyles such as the oval face shape, while others require the hairstyle to be more conscious about the face shape. Furthermore, visualizing a possible haircut before getting one can save a lot of frustration. Does a long, messy hairstyle fit better or maybe going blonde is a move.

Dataset

The dataset that I will be using for this project is a CelebA dataset, specifically, the dataset found here https://www.kaggle.com/kevinpatel04/celeba-original-wild-images#130007.jpg. This dataset contains over 202,599 face images each with 40 binary attributes that have been annotated. These attributes include hair color, facial hair, face shape, cheekbone structure, and more. These attributes are already used in the git project I will be working off, and I will need to use them even more to emphasize certain features. Furthermore, the git project has created training and dev/test sets from this dataset. The picture below shows the distribution of all attributes in the CelebA dataset. The blue means that attribute was present and the tan means that that attribute wasn't present.

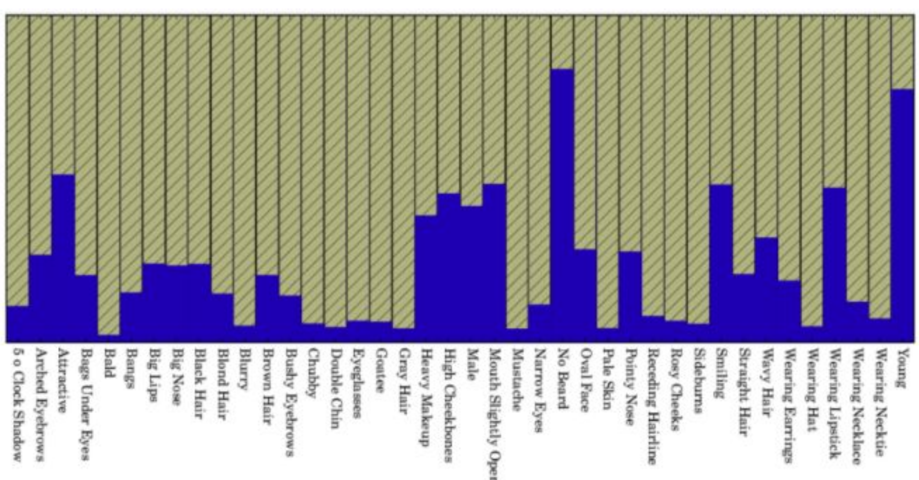

[1] https://www.thetrendspotter.net/find-the-perfect-hairstyle-haircut-for-your-face-shape/

Clearly there is some bias with young people and people with no beards. This means the neural network could underperform when inputted with images of older people, or people with beards. There is also bias towards individuals with Chubby, Eyeglasses and Bald attributes. So for the same reasons, inputs with those attributes might underpreform.

Hyperparameters and Architecture choices

Initially the program was going through each inner epoch loop at around 6 seconds. The GitHub repo had 900 inner epoch loops for each of the 200 epochs. This would take over 1,000,000 seconds per iteration. First I changed the batch size from 32 to 320. As Andrew Ng. described in lecture, finding a batch size close to one would be slow and equivalent to iterating over each image, while a batch size equal to the amount of images in the dataset would mean the program would lose the ability to learn after each iteration. However, when I increased my batch size to 320, my deep learning AWS instance crashed as it ran out of memory. So, I lowered the batch size to 96 which tripled my speed. It was still too slow so I reduced the number of epochs to 20.

I had a lot of trouble setting up the GitHub repository and getting it all to work with AWS so I didn't have much time to tune the hyperparameters.

I tuned the learning rate from .0002 to .0005 to .001. Below are the pictures of the loss visualized.

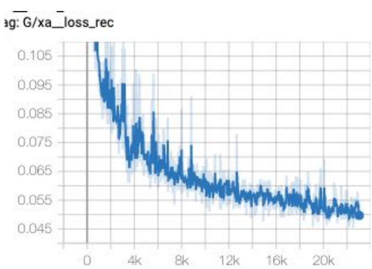

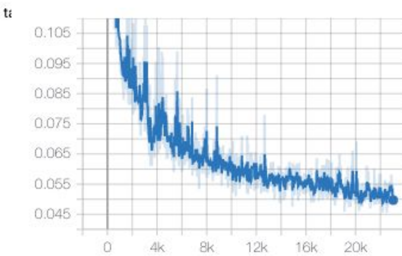

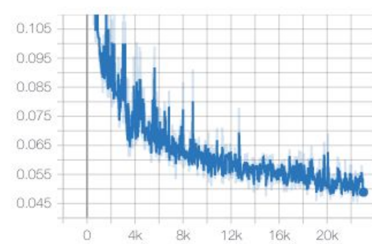

The loss for .0002 learning rate and the .0005 learning rate seem pretty minimal. The loss for the .001 learning rate has higher peaks and more fluctuation during the run when compared to the lower learning rates. However, the final loss value is slightly lower when compared to the lower learning rates. Also the loss for the .001 learning rate has a steeper function when compared to the lower rates, which means during earlier epochs, the program is learning faster.

I also played around with different optimizers, mainly Adam and RMSProp and ended up using the AdamOptimizer described in the lecture. This was also the recommended optimizer described in the STGAN paper.

Results

I tested the trained model with pictures from the CelebA dataset and professional headshots outside the dataset. I tested the Attributes Bald, Blonde Hair, Black Hair and Bangs and then compared the photos from the different distributions.

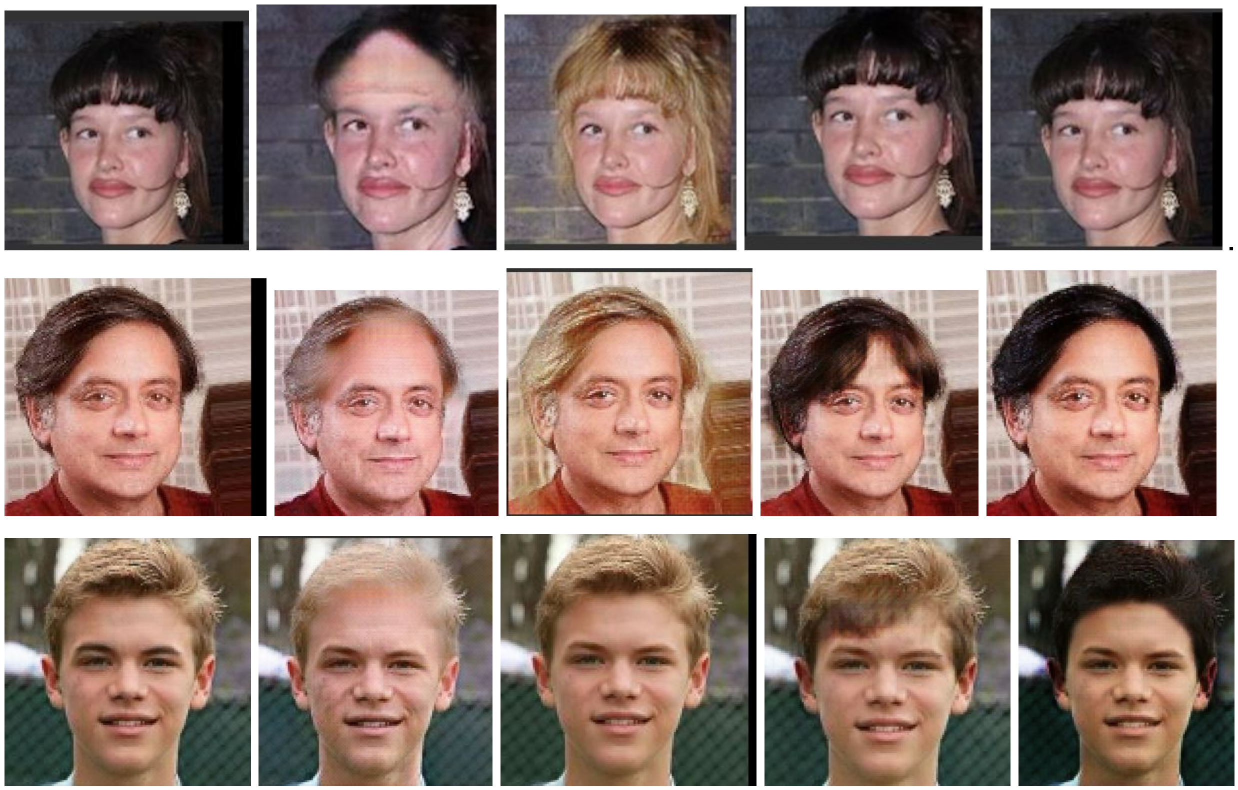

The order of the photos is the original then added attributes of bald, blonde, black hair and bangs.

Next I tested the photo on a headshot of me. This photo is not part of the training distribution but is high quality and straight on. The original photo was formatted to be 178 pixels by 218 pixels. The next photos are in the same order.

Result Analysis

Generally, this model worked well when transferring a hair to a different color. Adding a blonde or black hair attribute had very realistic results and definitely helped envision the change. The model struggled with changing hairstyles such as length and thickness. Adding bangs or removing hair had less than ideal results.

One interesting insight to note is when the model changes a face to bald, it often merges features associated with older individuals. Clearly, the model has added wrinkles to the forehead and the chin area in the photo below. It also slightly changed the hair color of the images to a slightly grayer version.

I believe this could be caused by two reasons. Firstly, the distribution in the CelebA dataset has few pictures of bald individuals who are younger. Since balding is often associated with aging, the model wouldn't be able to distinguish wrinkles from a property of aging rather than being bald. Another more plausible reason could be the possibility that the model wasn't trained long enough. This model was trained with 20 epochs, which was reduced from 200. As noted in lecture, training a model longer can reduce avoidable bias and get the model closer to Bayes optimal error.

Insights

This model, if surrounded by a friendly user interface could prove useful for many barbers, or individuals looking to change up their style. Companies who sell hair products could also utilize a model similar to this one to help customers envision hairstyles with the new hair product. This oftentimes happens to me as I wonder how a different hair gel would affect my hairstyle. This would make customers more informed about a certain product before buying it. My experience using this model has definitely helped me envision myself with a different hairstyle. Who knows, maybe I'll dye my hair blonde.