Gender Classifier and CycleGAN for Altering Facial Images

Srinivas Halembar, Mark Cramer, and Harry Jiang
{halembar, mdcramer, xinxuan} @stanford.edu

Abstract

We developed a gender classifier using transfer learning with a pre-trained VGG-16 model that achieved results comparable to state-of-the-art available gender classifiers. We also employed a CycleGAN to address the issue of gender conversion. We applied hyperparameter tuning, fine tuning and transfer learning to improve the accuracy of our models. We present the gender classifier performance along with sample images from gender conversion using a CycleGAN.

Dataset

We used the Diversity in Faces dataset released by IBM 4 weeks ago. The dataset consists of over 960 thousand images and are labeled for craniofacial features, skin color, age as well as **gender**, which is the feature we utilized.

Discussion & Future Word

Both the gender classifier and CycleGAN demonstrated amazing performance. In the future we'd like to train the classifier with even more hyperparameter turning and for a longer duration, while investigating better strategies to improve and evaluate GAN performance.

References

[1] Simonyan et all, Very Deep CNN for ... Image Recognition, 2014
[2] Zhu et al, Imaged-to-Image Translation … CycleGAN, 2017
[3] Merler et al, Diversity in Faces Dataset, Jan 2019

Models

For the gender classifier we utilized a CNN based on the VGG-16 model. We replaced the last FC layer of a pre-trained VGG-16 with multiple FC layers that performed binary gender classification. We trained it with 50% dropout on 855k images and then subsequently fine-tuned previous convolutional blocks of the VGG-16 model for even better results.

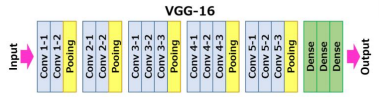

For our generative model we utilized an existing implementation of a CycleGAN. We defined the two domains to be *male* and *female*, respectively, and trained the GAN with the Diversity in Faces dataset.

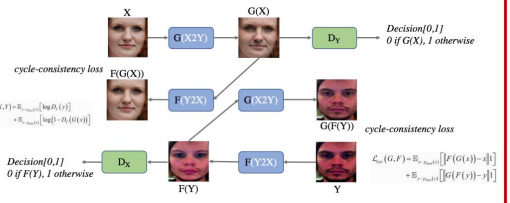

Results

The gender classifier eventually achieved an accuracy of 95.2%, with 2 FC layers and blocks 3, 4 & 5 of the original VGG-16 unfrozen. Below are some classification errors. We suspect many errors are mainly attributed to the dataset as well as the limit of the Bayes error.

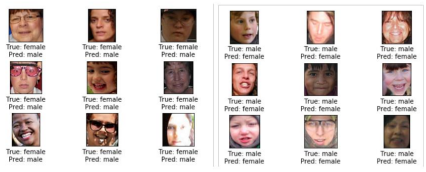

We experimented with two implementations of the CycleGAN, one with TensorFlow and one with Keras. We found that the TensorFlow model worked remarkably well and we noticed significant gender attribute changes before and after the CycleGAN's alterations. Some sample results are below – the first row are real images and the second row are generated.