Introduction

We first thought of exploring neural style transfer within the context of cartooning images; however, we realized that many projects have already explored this area and most neural style transfer projects focus on western styles. As three Chinese painting lovers, we decided to try out neural style transfer with the traditional Chinese ink painting style. We plan to exploit CNN for neural style transfer and GAN for image colorization. We will implement additional methods and constraints to make neural style transfer more compatible with Chinese painting techniques such as blank-leaving, brush stroke, or ink washing.

Challenges

We have modified ChipGAN and CycleGAN correspondingly to make them compatible with the dataset including images of ink-wash paintings of horses and landscapes. After we set up the environment correctly, although the model successfully started to train, the time required to complete the whole training process exceeded one day (one of our teammate's computers has equipped with the latest powerful GPU). Our model suffered from very low speed and high computation costs and we have not figured out an effective way to solve the efficiency problem. As a result, we had to cut down the size of the dataset and number of epochs for training for now before we confirmed our settings for the model are reasonable.

Insights

We greatly reduced the number of epochs trained, we have 30 epochs to train with the initial learning rate and another 30 epochs to train with a decaying learning rate. The results are actually better than we expected. We can see it transfer from a real phone to an ink-washing style painting.

Photos adapted from the third training epoch:

Plan

As I mentioned above, our primary concern is trying to reduce the training time for each epoch. Since we read from online studies that GAN benefits disproportionately from large mini-batch sizes, the batch size is now set up to 1, but a large number of batches are slow and expensive to emulate on conventional hardware such as our own laptops. It appears that altering the training parameters is insufficient to accelerate the training process. We will spend more time optimizing the structure of our model or looking up additional research paper that includes some useful methods to solve the efficiency problem of GAN-based approaches to neural style transfer. For example, we went through the idea of core-set selection from the paper Small-GAN: Speeding up GAN Training using Core-Sets, which offers an idea that draws a large batch of samples from the prior and then compresses that batch. We may generate a cached dataset of each training image's Inception activations, project them randomly to a lower dimension, and then apply core-set selection on those projected embeddings at training time to efficiently create large batches of real images. We are not sure if this technique will work on our dataset, but we would like to dedicate some time to research it. Moreover, our model now has quite a few overlapped structures and setting as the original design of ChipGAN and CycleGAN. As 2470 students, we are definitely going to explore something more innovative. Prior to experimenting with more radical modifications to the structure and settings of our model, we simply need a little more time to ensure that our current GAN model still generates results that are moderately dependable.