Visualizing adverse effects of Global Warming on Mountains through CycleGAN

Troy Shen (troyshen), Gunguk Kim (gunguk)

CS230 Deep Learning, Department of Computer Science, Stanford University

Stanford Computer Science

Abstract

The problem of neural style transfer has been approached in a variety of ways. We build on and modify Zhu et. al's CycleGAN with the purpose of finding a forward mapping function that maps from the domain of snowy mountain images to the domain of non-snowy mountain images. Because our motivations dictate that we generally care more about the effectiveness of the forward mapping function, we loosen the constraints on cycle loss consistency in the hopes of achieving more photo-realism. We hope our work sheds light on the ways that generative models can be improved upon for future research and in service of solving real-world problems.

Stanford University

Background

As a byproduct of climate change, the melting of mountains has become an increasingly relevant ecological problem. Our project involves the use of neural style transfer to create a visualization of how mountains would appear as a consequence of global warming.

Problem

Using CycleGANs, minimize cycle consistency loss and adversarial loss to find a forward mapping function that maps the domain of snowy mountains onto the domain of non-snowy mountains.

Data and Modifications

We used two datasets for this model: images of Yosemite during the summer and a hand-curated dataset of mountains reported to be melting at an alarming rate.

Our model builds on Zhu et. al's CycleGAN model through additional features: cycle consistency with weight decay, modification of the architecture through additional resnet blocks, and use of a skip connection.

Model and Methodology

The goal of the CycleGAN is to learn image-to-image translation from unpaired datasets. It achieves this through the use of two generators and two discriminators mapping images from the domain of snowy mountain images to the domain of non-snowy mountain images and vice versa.

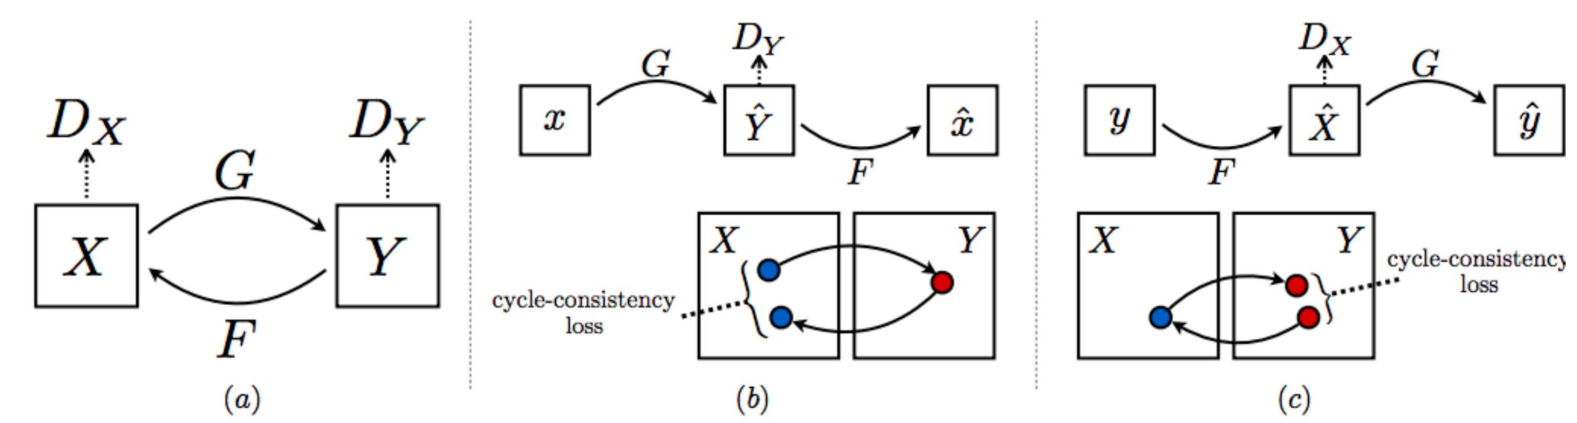

Figure 1: The CycleGAN Model

The overall objective (below) is to minimize the loss function, namely the adversarial loss of both generators and the cycle consistency loss. The cycle consistency loss aims to account for the fact that translating from one domain to the other and back again should result in the original image.

$$\mathcal{L}(G,F,D_X,D_Y) = \mathcal{L}_{GAN}(G,D_Y,X,Y) + \mathcal{L}_{GAN}(F,D_X,Y,X) + \lambda\mathcal{L}_{cyc}(G,F)$$

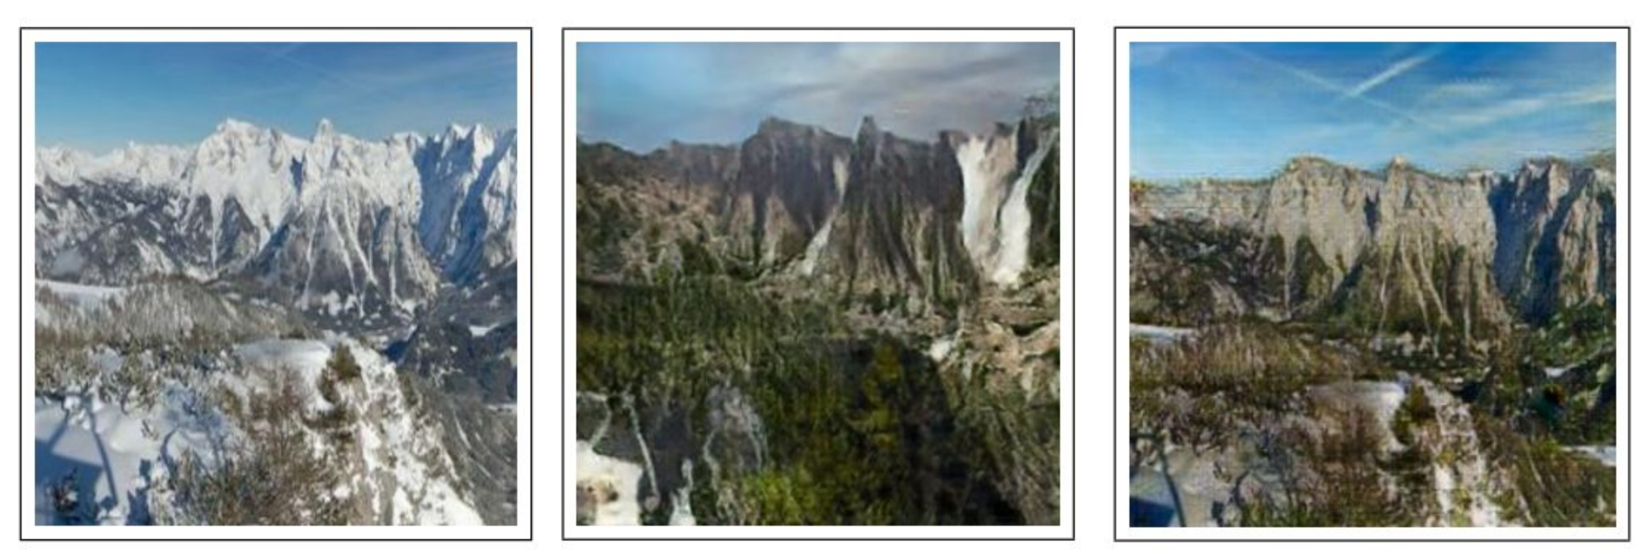

(a) Input Image (b) Without Architectural Change (c) With Architectural Change

Figure 7: Comparison between images generated with and without architectural.

Results

$$\mathcal{L}(G,F,D_X,D_Y,t) = \mathcal{L}_{GAN}(G,D_Y,X,Y) + \mathcal{L}_{GAN}(F,D_X,Y,X) + \lambda_t\mathcal{L}_{cyc}(G,F)$$

(a) Input Image (b) Generated w/o Weight Decay (c) Generated w/ Weight Decay

Figure 6: Comparison between images generated with and without weight decay.

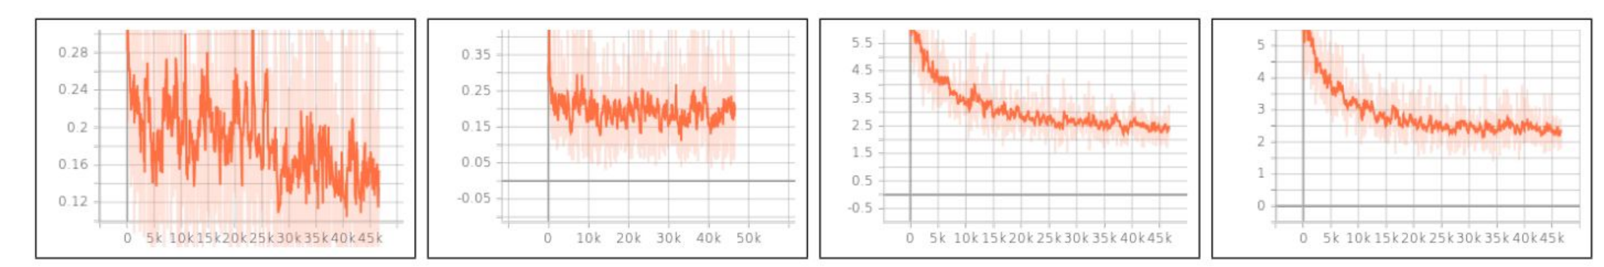

(a) Discriminator A Loss (b) Discriminator B Loss (c) Generator A Loss (d) Generator B Loss

Figure 3: Initial Loss Functions

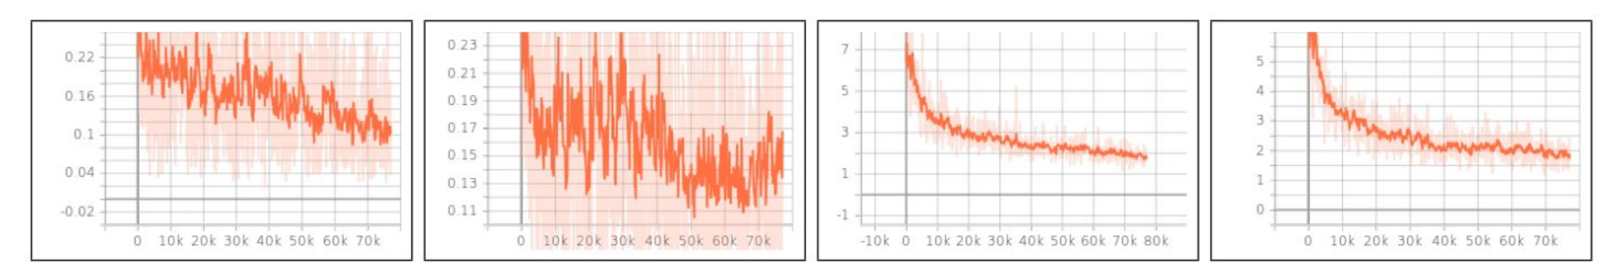

(a) Discriminator A Loss (b) Discriminator B Loss (c) Generator A Loss (d) Generator B Loss

Figure 4: Loss Functions with Weight Decay

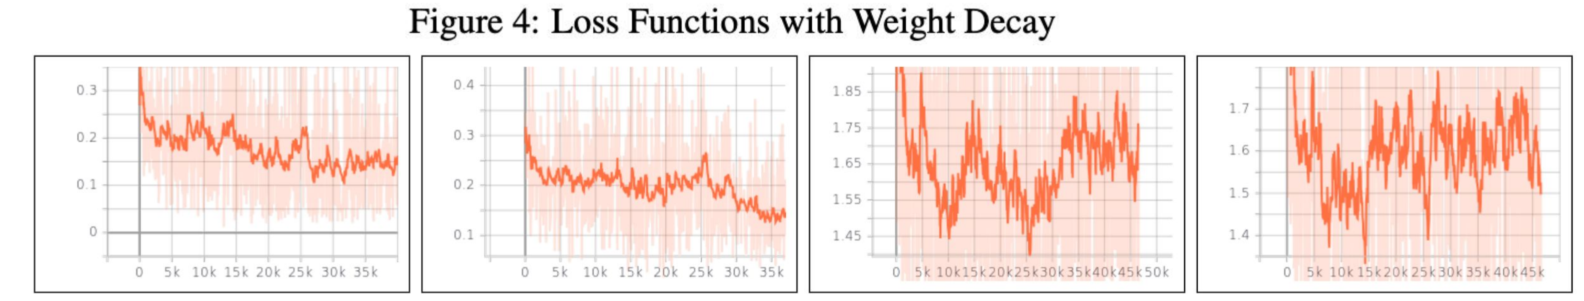

(a) Discriminator A Loss (b) Discriminator B Loss (c) Generator A Loss (d) Generator B Loss

Figure 5: Loss Functions with Architectural Changes and Weight Decay

Analysis and Future

Although our model was mostly successful, modifications that could be explored include using larger datasets, tuning of hyperparameters, and other strategies like: latent generative models, flow normalization to minimize cycle consistency loss, and a HarmonicGAN implementation.