Generative Art via Neural Networks

Evan Russenberger-Rosica
er1093a@student.american.edu

Department of Mathematics and Statistics - American University

Learn more at evanrosica.com/GANArt/

Abstract

- We investigate if neural networks can be used to create art which is novel, detailed, and aesthetically pleasing. We evaluate two state-of-the-art approaches against these criteria: neural style transfer (NST) as in Johnson et al. and unpaired image-to-image translation (IIT) as in Zhu et al.
- In the process, we mathematically prove that Zhu et al's implementation of the cycle consistency constraint is unnecessarily complex. We also observe that their method undesirably converges to the test image as the test image becomes relatively large.
- We conclude that both methods meet our criteria in different ways. NST works best when trained on abstract images as well those with geometric patterns, while IIT works better on less abstract images.
- We believe that our comparison between IIT and NST is the most extensive yet done, and the only one to consider very large images.

Unpaired Image-To-Image Translation

- In *Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks*, authors Zhu et al. define **image-to image translation** as "a class of vision and graphics problems where the goal is to learn the mapping between an input image and an output image using aligned image pairs."
- Since paired image training data is rare, Zhu et al. investigate how to translate images from a source domain X to a target domain Y in the absence of paired examples.[zhu, 1]
- Zhu et al improve on standard GAN optimization problem by adopting the bias that translation should be **cycle consistent**. I.e., if we were to translate a sentence from English to French, and then translate it back from French to English, we should arrive back at the original sentence.

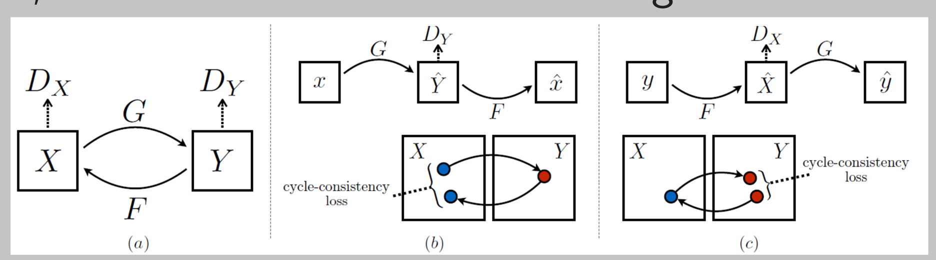

- Mathematically, if $G : X \to Y$ and $F : Y \to X$ then
$$F(G(x)) = G(F(x)) = I(x) = x$$
where $I(x)$ represents the identity function. Zhu et al. state "G and F should be inverses of each other, and both should be bijections."[zhu, 2]
- In the following theorem, we prove that this is actually a stronger condition than is needed to ensure cycle consistency.

Theorem

If $F(G(x)) = I(x)$ and F is an injection, then $G(F(x)) = I$.
Proof: By contradiction, assume $F(G(x)) = I(x)$ and $G(F(x)) \neq I$. Then $\exists x \in Y, G(F(x)) = y \wedge y \neq x$. By hypothesis we know $F(G(x)) = I(x)$, therefore:

$$\begin{aligned} G(F(x)) &= y \\ F[G(F(x))] &= F(y) \\ I(F(x)) &= F(y) \\ F(x) &= F(y) \end{aligned}$$

which is a contradiction since F is an injection. Therefore $G(F(x)) = I$ and so F, G are cycle consistent. $\square$

Neural Style Transfer

- Unlike CycleGAN which learns to mimic the style of an entire collection of artworks, "**neural style transfer** learns to transfer the style of a single selected piece of art onto another".[johnson][gatys]
- The neural style transfer system consists of two parts: an image transformation network, and a loss network. "The image transformation network is a deep residual convolutional neural network. By jointly minimizing the feature reconstruction loss and a style reconstruction loss also based on features extracted from a pretrained convolutional network..... this method produces high-quality results."

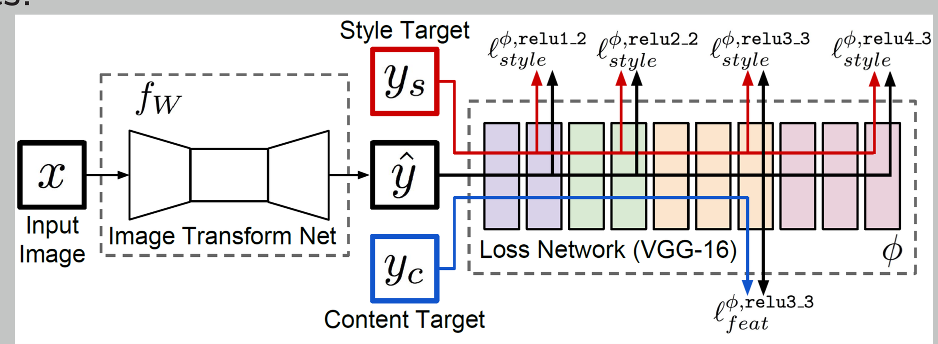

Training and Logistics

- Since NST uses a pre-trained neural network, the user need only supply a style image and a content image in order to produce results with NST.
- A style transfer network can be trained on a modern GPU in several hours. (We used an AWS p3.2xlarge instance with a Tesla V100 GPU w/16GB VRAM). Once the network is trained, performing the style transfer itself takes only a few seconds.
- Conversely, to use CycleGAN, one must first build a (large) image dataset, and then train the GAN. This is a very resource intensive process, and is best done on a cloud service providing GPU compute such as AWS.

Neural Style Transfer

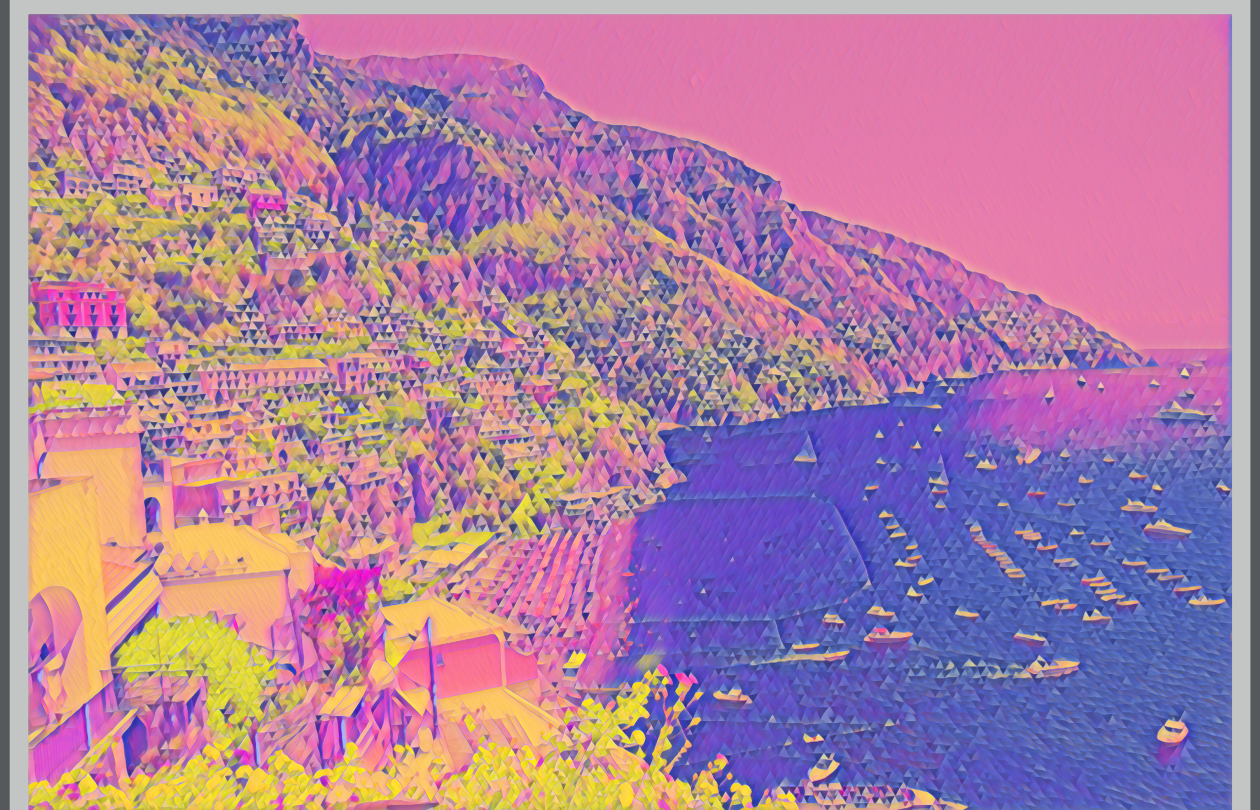

Neural Style Transfer - Style Images

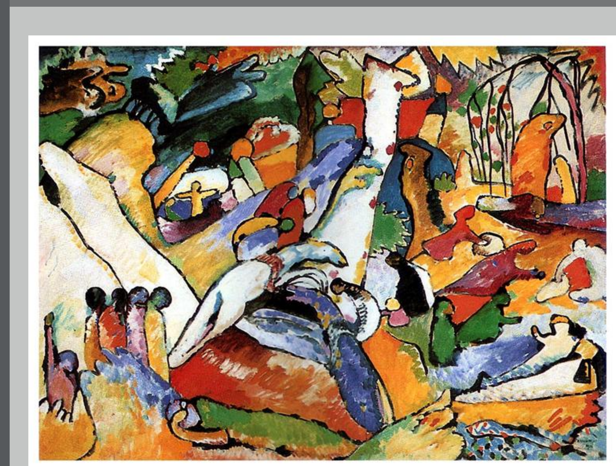

Figure: Kandinsky: Composition II, 1910

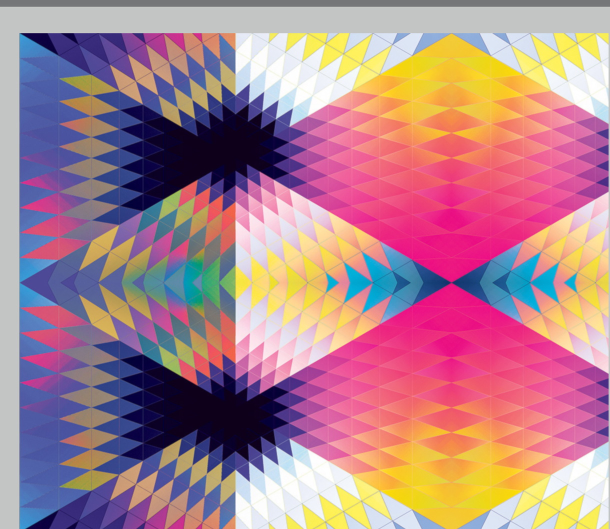

Figure: Unknown

- The network which created the top image in this column, seems to have learned to make brushstrokes from Kandinsky's Composition II (above left).
- Close inspection of the second image from the top shows that the style network has learned to render objects as triangles. The unnamed style image is above right.

CycleGAN: Monet, Ukiyo & Cezanne Styles

Reference Images

- We use the above original pictures of the Almafi Coast of Italy as reference content images.