Neural Style Transfer for Understanding How Stylistic Choices Affect Emotional Affect

Alex Weitzman
MS in Computer Science
alexwe@stanford.edu

Phoebe Yao
BS in Human-Centered Design Engineering
phoebeyao@gmail.com

Predicting

We used Neural Style Transfer on landscape paintings to qualitatively analyze which stylistic choices in a painting can give the painting a positive or negative emotional effect for viewers. We grouped landscape paintings into a valence category ("positive" or"negative") and used this group as style reference images for Neural Style Transferto try to change the emotional affect of a neutral content landscape painting. We also explored optimization methods such as color preservation and photorealism regularization to improve the appearance of results, and we ran a qualitative survey to investigate the performance of our algorithm.

Features

We experimented with the following hyperparameters:
Average pooling: We modified the standard max-pooling layers in order to improve the gradient flow in the CNN. Average pooling resulted in more aesthetic and natural appearing images. Since max pooling takes the max value among pixel values, generated images typically result in more salient features.

Content-Style Trade-off: We varied the ratio between the emphasis on the content and style loss (α/β) as an extension to the loss function. We found a ratio of 5e-2 content to style worked the best on landscape paintings. This also appears to be the standard ratio in NST literature.

Discussion

We noticed interesting trends in the generated paintings: the positive valence images we generated were more colorful, glowing, and detailed than the negative valence paintings. When we surveyed 22 participants what valence they interpreted from the paintings, many of the responses were what we expected, but some remained controversial. After all, art depends on the viewers' knowledge, symbolism, artist

Data

Our original dataset included 2813 images scraped from Wikiart.org that were tagged with 10 discrete emotions: joy, optimism, fear, surprise, lust, anger, sadness, neutral, envy, and love. We received the dataset from our mentor, Chez Mana, who is sponsoring the project idea with the approval of CS230TAs. The people who tagged the images both have university degrees, one an artist and the other a computer programmer. They concentrated on the primary emotion of every artwork, and we later altered this to valence: positive, negative, neutral.

Results

Positive Valence Generated Images

iteration: 1000
content loss: 290151
style loss: 1.44825e+06

iteration: 5000
content loss: 330623
style loss: 1.45498e+06

Future

We would love to look into Generative Adversarial Networks and tackle the challenge of classifying the emotion of a painting based on content, using that as a discriminator while we generate different paintings to try to teach a model how to create paintings that capture a certain valence or emotion.

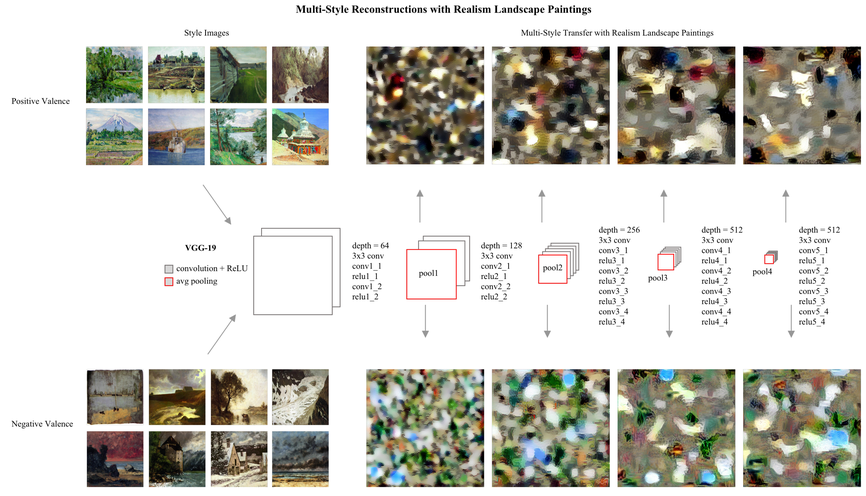

Models

We used the VGG-19 network trained on the ImageNet database. Per standard NST models, our training focuses on minimizing the following loss function:

$$L = \left|\left|Content_C - Content_G\right|\right|_2^2 + \left|\left|Style_S - Style_G\right|\right|_2^2$$

This function aims to minimize the difference in content between the generated image and the content image, as well as the difference in style between the generated image and the reference style image. In our case, we augmented the above loss function to train our Multi-Style NST by taking the sum of the loss over n-chosen style images.

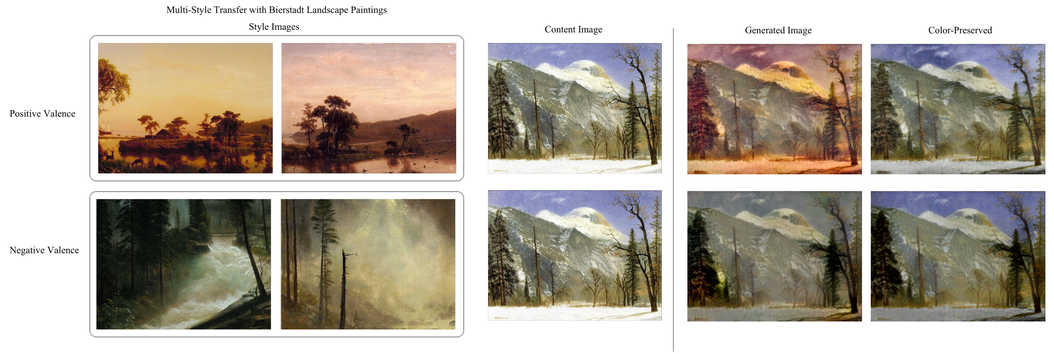

References
[1] anishathalye, Neural-Style Implementation. https://github.com/anishathalye/neural-style
[2] TensorFlow 1.13.1
[3] L. A. Gatys, A. S. Ecker, and M. Bethge. A neural algorithm of artistic style. *arXiv preprint arXiv:1508.06576,* 2015.
[4] L. A. Gatys, A. S. Ecker, and M. Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016.
[5] L. A. Gatys, M. Bethge, A. Hertzmann, and E. Shechtman. Preserving color in neural artistic style transfer. *arXiv preprint arXiv:1606.05897*, 2016.
[6] F. Luan, S. Paris, E. Shechtman, and K. Bala. Deep photo style transfer. In *CVPR*, 2017.
[7] Daan Wynen, Cordelia Schmid, and Julien Mairal. Unsupervised learning of artistic styles with archetypal style analysis. *arXiv preprint arXiv:1805.11155*, 2018.
[8] V. Dumoulin, J. Shlens, and M. Kudlur. A learned representation for artistic style. In *ICLR*, 2017.
[9] X. Huang and S. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. *arXiv preprint arXiv:1703.06868*, 2017.