Mama?



Sophia Kolak

DESCRIPTION

Concept

This project was inspired by the shifting relationship between content and form in art history. Whereas renaissance art typically depicted imaginary figures with realism, early modern art depicted real people abstractly. How are we to make sense of this change---of the shifting balance between abstraction in art and the contemporary experience of viewing it? To explore the relationship between content and form across movements in art, we choose two images depicting mother and child from different time periods and artistic movements. We then use machine learning (style transfer) to create new images that combine the popular style of one era with the popular content of another.

The first image, Michelangelo's Pieta, embodies renaissance style art. It portrays the Virgin Mary cradling Jesus, with striking mastery of the human form. Notwithstanding the disproportionate size of Mary (a pragmatic choice), the bodies are depicted with such anatomical precision that they appear to be real people frozen in a moment. It is difficult to imagine that the medium of Mary's skirt is in fact marble and not the smooth cloth Michelangelo has tricked us into seeing. And yet, when we consider the subject of this art, all realism vanishes. It is a visual depiction of an imaginary scene, a representation of iconic figures in an idealized religious moment. The aesthetic value in this work is tied to its content, namely, the glorification of Christian imagery through technical mastery. Such value exists because Michelangelo could make real what was imagined, he could give form to ideas, adding shape and structure to something ephemeral, by projecting a cherished legend onto natural, human bodies.

Conversely, the second image, Monet's "Mother and Child", instead depicts a real subject in a loose, impressionistic form. Unlike renaissance art, which featured naturalistic depictions of bodies in practical earthly forms, impressionism depicts real subjects in an abstract form. Known for its rapid, un-calculated brush strokes, impressionism forces us to reconsider what qualities make art masterful. Prior to this movement, art was largely evaluated based on how closely it was able to mimic reality. Beauty was in capturing, not interpreting. Impressionism is an abandonment of this pure formalism, an attempt to represent the fleeting moment in whatever way our limited human hands and minds can. Interestingly, with slightly less realism, both the viewer and the artist become more involved with the art. Now the artist has to decide which parts of the picture to fill in. Similarly, the viewer has to read an image that is purposely left partially blank, leaving holes for their imagination and subjective perspective on the scene.

By using machine learning to transfer the impressionistic style of "Mother and Child" to the religious content and masterful form of "Pieta", we can explore the synthesis of renaissance and impressionistic beauty. Additionally, by choosing two images which both centrally feature a mother and child, we keep the underlying theme of each image constant, creating a thought provoking final result.

Technique

To create this image, I used the neural-style algorithm (implemented in Pytorch) to transfer the style of one image onto the content of another. As suggested in the original paper, I used the L-BFGS algorithm to run gradient descent. Also following the paper, I used a pre-trained 19 layer VGG network. For image generation, I used 500 steps, saved the image after every 20 steps, and picked the one I thought looked most visually interesting.

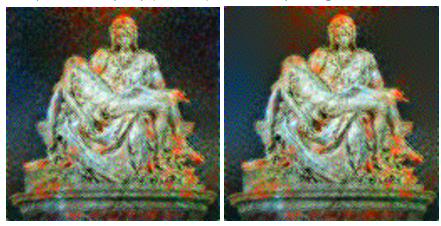
Process

My process mainly involved tweaking various parameters to achieve an interesting balance of style and content. Originally, I had the style weight set to 100,000,000 and the content weight set to 1. I found that this produced images in which the original image was too difficult to make out, so I reduced the style weight to 1x10^8 and then to 1x10^6. I also experimented with a varying number of steps, ranging from 200-800. I found that the most interesting images were produced around 200-300 steps, but kept the steps set to 800 and simply printed the output of the image at every 20 steps so I would have more options. The final image is the result after step 800 with style weight=1x10^9

Example from when the style weight was 1 x 10^8... quite like these



Example from early step (100/140) when the style weight was 1x10^6, these look to digital IMO



I also tried transferring the images in the reverse order (i.e., transferring the style of Pieta to Mother and Child), but I found the results from this transfer to be much less aesthetically pleasing, and honestly a bit creepy.

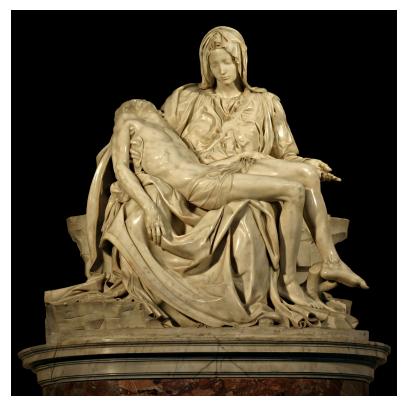


Reflection

As mentioned above, I chose the final image mainly just based on my own intuition after tweaking various hyper-parameters of the model, particularly looking for a configuration that produced a decent balance between style and content loss. I was pretty satisfied with the final result. To me, it shows what it would have looked like if Michelangelo was an impressionist instead of a renaissance painter/sculptor.

RESULT

The original images used for this style transfer:



Michelangelo, Pieta 1499



Monet, Mother and Child 1875

CODE

Repository hosting my code/Jupyter notebook I used for this project -> https://github.com/sophiakolak/art-ml-2

The reference off of which my code is largely based -> https://pytorch.org/tutorials/advanced/neural-style-tutorial.html

REFERENCE

Style transfer paper: https://arxiv.org/abs/1508.06576