

In-Image Style Blending

Tom Henighan

Stanford Department of Physics

henighan@stanford.edu



Neural Style Transfer Is a method for transferring the “style” of one image onto another using Conv Nets[1]

Algorithm: The “style” image and “content” image are both fed through the convolutional layers of a conv-net pre-trained for some visual task (e.g. recognition).

Use gradient ascent to make a new image with features matching those of the content images in a spatially resolved way (match conv layer outputs) and matching those of the style images in a spatially un-resolved way (matching Gram matrices)


Gram Matrix




H x W x D output from conv-net layer

1 Flatten:
2 HW x D
Call F


$$G_{ij} = \sum_k F_{ki} F_{kj}$$




Masking: Allow style transfer for only part of the image [2]





How to determine mask for each Layer?



Multiple Style: Use multiple masks with multiple styles



Preserve Colors: By using only luminance channel of style and content images can somewhat preserve color [2]



[1] 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.

[2] L. A. Gatys, A. S. Ecker, M. Bethge, A. Hertzmann, and E. Shechtman. Controlling perceptual factors in neural style transfer. arXiv preprint arXiv:1611.07865, 2016.