



Segmented Neural Style Transfer for Chair Generation

Sam Premutico

samprem@stanford.edu - December 2018

Overview

- We attempt to augment neural style transfer with object localization and image segmentation in order to style a *portion* of the content image in the style of a *portion* of the style image
- In particular, we build off previous localized-style transfer work to generate novel *chair* designs by first identifying the portion of the content and style images that contain the respective chairs, and second styling the region of the content image in the style of the style image
- As generic neural style transfer styles the *entire* content of the content image in the style of the *entire* style image, it is not effective at generating subtle styling of particular regions of the content image in the style of a particular region of the style image

Dataset

- Collected ~100 color images from West Elm, Ikea, and Design Within Reach websites
- Collected images both with white background as well as with full background, i.e. on display in a fully furnished room
- Trimmed each to 255 pixel maximum height/width
- Generate bounding boxes and segmentation masks for images using GrabCut segmentation algorithm^[1]




Model

- Segment images using *GrabCut* algorithm
- Feed style, content images to pre-trained VGG-19 network along with segmentation masks for style image(s) ^[3]
 - 16 Convolutional → ReLU layers & 5 Pooling layers
- Calculate style loss with segmentation-aware loss function


Style Loss ^[2]

$$\frac{\sum_{i,j} (G_{ij} - A_{ij})^2}{4N^2M^2}$$



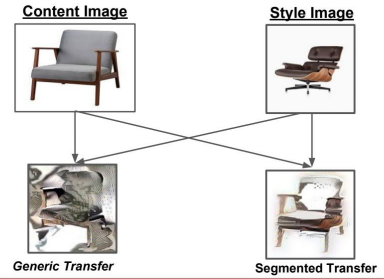
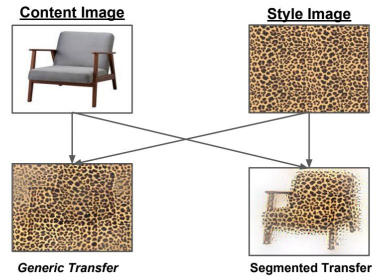
- To calculate *style loss* at layer l , we compare the Gram matrices of generated image G and original style image A at layer l
- We isolate the portion of the content image that we want to style by performing an element-wise multiplication of G and A with our segmentation mask B transformed at layer l

Segmentation Mask Across Network



Results

(so far)



Discussion and Next Steps

- While we were able to apply localized-style transfer to a region of the content image, we've struggled to transfer the style of a *portion* of the style image onto the the content image
- We hope to develop an effective method of segmenting the style image and extracting the style from a particular object in the style image
- Further tuning of the layers at which we compute style loss may improve the quality of the generated images

1. Rother et al., "GrabCut: Interactive Foreground Extraction Using Iterated Graph Cuts"
 2. Gatys et al., "Image Style Transfer Using Convolutional Neural Networks"
 3. Segmented Transfer: <https://github.com/cysmith/neural-style-tf>