

Exploring the Effectiveness of Mask-Guided Feature Modulation as a Mechanism for Localized Style Editing of Real Images (Student Abstract)

Snehal Singh Tomar, Maitreya Suin, A.N. Rajagopalan

Indian Institute of Technology Madras
 snehal@smail.iitm.ac.in, maitreyasuin21@gmail.com, raju@ee.iitm.ac.in

Abstract

The success of Deep Generative Models at high-resolution image generation has led to their extensive utilization for style editing of real images. Most existing methods work on the principle of inverting real images onto their latent space, followed by determining controllable directions. Both inversion of real images and determination of controllable latent directions are computationally expensive operations. Moreover, the determination of controllable latent directions requires additional human supervision. This work aims to explore the efficacy of mask-guided feature modulation in the latent space of a Deep Generative Model as a solution to these bottlenecks. To this end, we present the SemanticStyle Autoencoder (SSAE), a deep Generative Autoencoder model that leverages semantic mask-guided latent space manipulation for highly localized photorealistic style editing of real images. We present qualitative and quantitative results for the same and their analysis. This work shall serve as a guiding primer for future work.

Introduction

The Swapping Autoencoder (SAE) (Park et al. 2020) is a Deep Generative AE Model which learns separate latent representations for the structure (S_s) and style (S_t) information present in images. We exploit the Modulated Convolutional layer of the SAE, where the global style vector is spatially broadcasted and multiplied with the structure tensor S_t for modulating the style of the generated image. We claim that, instead of repeating the same style vector for all pixel locations, if we can first localize a particular Region of Interest (ROI) and then manipulate the style content by adding noise in the latent space for only that region, it will result in a locally edited image. The resultant image often depicts coarse style edits on a ROI. To refine this output, we employ a refinement block that acts on the coarsely edited ROI while being guided by a discriminator to generate realistic styles. Our framework does not require any ground truth for the locally edited images and is only weakly supervised by GT segmentation masks. Since our method does not require the identification of controllable latent directions post training, we eliminate the additional human supervision required by SOTA methods to achieve similar results.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

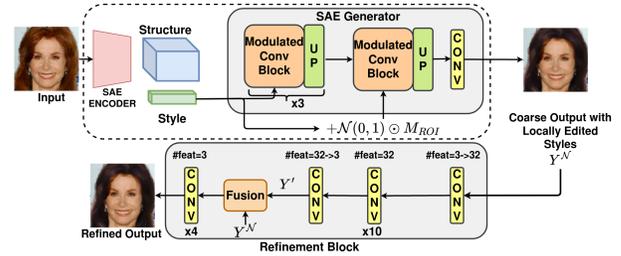


Figure 1: An illustration of the SMN (enclosed in ---) and the Refinement Block. M_{ROI} represents the SMPN’s output.

Method

Fig. 1 depicts our editing pipeline which consists of the Style Manipulation Network (SMN) and the Refinement Block. The SMN makes use of masks generated by the Semantic Mask Prediction Network (SMPN). Given an image, the SMPN first generates segmentation maps of different ROIs. Guided by these segmentation maps, the SMN performs local style modification in the input image.

Semantic Mask Prediction Network (SMPN) We build a semantic binary mask prediction network to localize individual regions of interest, namely: hair, skin, nose, eyes, and (lips + mouth). An overview of our SMPN is shown in Fig. 2. We employ a pretrained ResNet18-based framework for the encoder with four levels. Further, we introduce a few parallel convolutional blocks at every level. In the decoder, we use two convolutional layers at each level, followed by an upsampling operation. In the end, we use a sigmoid operation to produce a spatial mask and train using binary cross-entropy loss with the ground-truth binary mask. We use five SMPNs with the same structure for five ROIs.

Style Manipulation Network (SMN) We use a pretrained SAE (Park et al. 2020) to build our SMN. Given an input images of size $H \times W \times 3$, the encoder generates a structure tensor (S_s), having dimensions $H/16 \times W/16 \times 8$ and texture vector (S_t), having dimensions $1 \times 1 \times 2048$. The generator starts with S_s and gradually upsamples and refines it to finally reconstruct the input image. The key building block of the SAE generator (decoder) is the Modulated Convolutional layer. We argue that, as the same style vector (S_t)

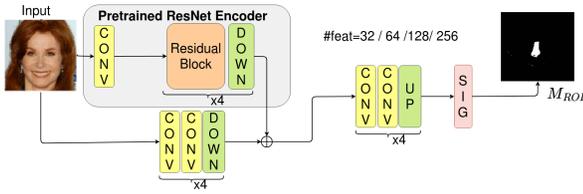


Figure 2: An overview of the Semantic Mask Prediction Network. Five networks with the same structure are used for predicting five different ROIs.

is repeated (to form S_t^{spat}) and applied to all the pixel locations, it results in uniform style changes in the resultant image. Thus, if we vary the style vector in the spatial domain, we will be able to manipulate the local style of any particular region. Driven by this, we deploy a semantic-mask-guided style vector modulation technique in our work. Given S_t^{spat} and a mask M_{ROI} indicating a particular ROI, we perform the following operation

$$S_t^{spat} = S_t^{spat} + \mathcal{N}(0, 1) \odot M_{ROI} \quad (1)$$

where $\mathcal{N}(0, 1)$ denotes the standard Gaussian Noise. Thus, we add ROI-specific noise to the latent style/texture vector S_t^{spat} , which changes the style content only for the corresponding ROI. As indicated in Fig. 1, we empirically selected the last but one decoder layer to add noise and found the changes to be realistic.

Refinement Block (RB) We refine our results using the RB. We use skip connection-based convolutional layers for the RB that operates at the full spatial resolution to preserve the finer pixel details. Let the output of the SAE after adding spatial-noise to S_t be Y_{SAE}^N . We apply a few convolutional operations on Y_{SAE}^N to generate Y_{SAE}' . Next, we fuse the refined output with the original output Y_{SAE}^N as follows

$$Y_{SAE}'' = Y_{SAE}' \odot M_{ROI} + Y_{SAE}^N \odot (1 - M_{ROI}) \quad (2)$$

We further pass Y_{SAE}'' through a few convolutional layers to generate the final refined output Y_{ref} . In our work, we use five separate refinement blocks for five ROIs.

Objective Functions The Semantic Mask Prediction Network (SMPN) is trained using standard binary cross-entropy loss

$$\mathcal{L}_{SMPN} = \sum_{ROI} (M_{ROI}, M_{ROI}^{GT}) \quad (3)$$

The SAE is pretrained with standard following the standard training procedure described in (Park et al. 2020)

$$\begin{aligned} \mathcal{L}(Y_{SAE}, Y) &= \mathcal{L}_{rec}(Y_{SAE}, Y) + 0.5L_{GAN, rec}(Y_{SAE}, Y) \\ &+ 0.5L_{GAN, swap}(Y_{SAE}, Y) + 0.5L_{CooccurGAN}(Y_{SAE}, Y) \end{aligned} \quad (4)$$

Here, Y_{SAE} refers to the reconstruction obtained from the decoder for the latent space $\{S_s, S_t\}$ without any noise, Y refers to the ground-truth (input) image. Other notations in Eq. 4 have been used exactly as in (Park et al. 2020). We train our refinement blocks to reconstruct the region corresponding to $1 - M_{ROI}$ same as Y_{SAE}^N . The style manipulated M_{ROI} region is trained with standard discriminator

Method	FID ↓	LPIPS ↓	Time Taken (s) ↓
SemanticStyleGAN	-	-	120.602
Ours	9.83255	0.1252	0.01143

Table 1: Quantitative results for our method’s performance versus SemanticStyleGAN (Shi et al. 2022). We evaluate the perceptual similarity of edits obtained (FID, LPIPS) per ROI with respect to the corresponding input and the time taken to perform one texture edit per ROI.

loss. The overall loss function can be expressed as

$$\begin{aligned} \mathcal{L}(Y_{ref}, Y_{SAE}) &= \sum_{ROI} (\mathcal{L}_{rec}(Y_{ref} \odot (1 - M_{ROI})), \\ &Y_{SAE} \odot (1 - M_{ROI}) + 0.5L_{GAN, rec}(Y_{ref}, Y_{SAE})) \end{aligned} \quad (5)$$

Results and Conclusion

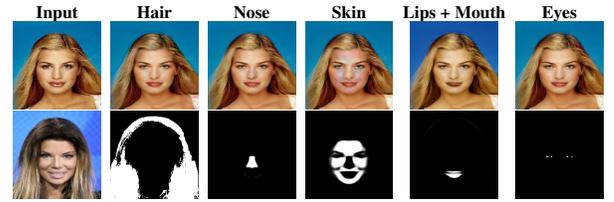


Figure 3: Row 1 - Localized Style Editing results. Semantic ROIs of the input image have been edited in accordance with the column label. Row 2 - SMPN output.

Figure 3 and Table 1 present our model’s qualitative and quantitative results, respectively. The perceptual metrics and qualitative results illustrate that our method performs ROI-wise localized, photorealistic, and structure preserving edits on real face images. The comparative analysis with SOTA in localized real image editing (Shi et al. 2022) using computation time reveals the massive speed-up, our approach offers over latent-inversion based SOTA methods. Moreover, our SMPN also produces highly accurate binary semantic masks for the intended ROIs, given an input image. Dataset, implementation, and training details have been provided in the supplementary material. In conclusion, We observe that although the results fulfill the intended objectives (clearly visible upon close inspection), the style effects are not very highly pronounced. We attribute this to the excessive refinement caused by the RB. As a future direction, we shall study ablations of the RB that yield best results and improve its architecture.

References

Park, T.; Zhu, J.-Y.; Wang, O.; Lu, J.; Shechtman, E.; Efros, A. A.; and Zhang, R. 2020. Swapping Autoencoder for Deep Image Manipulation. In *Advances in Neural Information Processing Systems*.

Shi, Y.; Yang, X.; Wan, Y.; and Shen, X. 2022. SemanticStyleGAN: Learning Compositional Generative Priors for Controllable Image Synthesis and Editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 11254–11264.